

Politecnico di Torino

ScuDo
Scuola di Dottorato ~ Doctoral School
WHAT YOU ARE, TAKES YOU FAR

Doctoral Dissertation

Doctoral Program in Electrical, Electronics and Communications Engineering
(38^{th}cycle)

Electronic Instrumentation for Biological Signal Acquisition

From Neural Interfaces to Cardiovascular Health Monitoring

By

Marco Boscherini

******

Supervisor(s):
Prof. Danilo Demarchi, Supervisor
Dr. Alessandro Sanginario, Co-Supervisor

Doctoral Examination Committee:
Dr. Benjamin Metcalfe, Referee, University of Bath
Dr. Dante G. Muratore, Referee, Delft University of Technology
Prof. Fabrizio Riente, Committee Member, Politecnico di Torino
Dr. Yuri Ardesi, Committee Member, Politecnico di Torino
Dr. Guido Pagana, Committee Member, Politecnico di Torino

Politecnico di Torino
2026

Declaration

I hereby declare that, the contents and organization of this dissertation constitute my own original work and does not compromise in any way the rights of third parties, including those relating to the security of personal data.

Marco Boscherini
2026

* This dissertation is presented in partial fulfillment of the requirements for **Ph.D. degree** in the Graduate School of Politecnico di Torino (ScuDo).

I would like to dedicate this thesis to Ellie, my life partner, who supported me unconditionally and gave me the right balance to overcome every obstacle encountered during the last four years.
To my parents and sisters, who never stopped believing in me, even from afar.
To my friends and colleagues, with whom I shared every single difficulty.
To my supervisors, Alessandro and Danilo, who gave me the opportunity to undertake this tough but extraordinary experience.

Abstract

The motivation of this thesis is to translate advances in sensing, electronics, and embedded processing into biomedical technologies that are more accessible, efficient, and responsive to real needs. This requires complete instruments to be designed around the biological information to be measured rather than by optimizing isolated components. An application-driven methodology is therefore applied to multichannel extracellular recording from *in vitro* neuronal cultures and non-invasive arterial pulse-wave monitoring.

The main contribution is an accessible event-driven neural acquisition platform for passive Microelectrode Arrays (MEAs). Extracellular spikes are detected in hardware and represented through channel identity and millisecond-scale timing, avoiding continuous sampling when waveform morphology is not required. The platform is modular and comprises a low-noise analog front end, a motherboard with the threshold-based event-driven core and two versions of control units: a 15-channel microcontroller unit, used as proof of concept, and a 60-channel Field-Programmable Gate Array (FPGA) unit, used as the ultimate implementation. Preliminary tests demonstrated feasibility of microvolt-scale signal acquisition with an input-referred RMS noise below 10 µV. At $\mathrm{SNR}_{\mathrm{lin}} = 10$, the MCU-based version achieved an aggregate F_1 score of approximately 0.98 up to 300 Hz, while the FPGA-based implementation maintained F_1 values of approximately 0.95–0.97 up to 200 Hz. Thanks to the proposed approach, the ideal continuous sampling with 12-bit payload at $10\,\mathrm{kS\,s^{-1}}$ was reduced by 99.1% for 15 channels and 98.7% for 60 channels. Mean system latency ranged from 0.34 ms to 1.45 ms. The complete prototype had an estimated hardware cost of approximately USD 3,000, excluding development and personnel costs. Its main trade-off is the loss of waveform morphology, while complete end-to-end biological validation remains necessary future work.

The second contribution was a pulse-wave monitoring platform integrating a custom rGO-R1 piezoresistive sensing material with a flexible PCB. A mechanical interface and embedded acquisition electronics were developed around the sensor to assess its feasibility for acquiring superficial arterial pressure pulses. Sensor-level characterization showed a low-pressure sensitivity of $0.281\,\mathrm{kPa}^{-1}$ and favourable short-term cyclic stability. The platform also acquired single-site and simultaneous carotid–femoral pulse waves. Although the physiological evaluation was preliminary, the identified structural and operating limitations guided the development of the subsequent wearable system.

The third contribution followed the rGO-R1-based sensing platform and addressed its limitations by adopting more stable and accessible commercial Force-Sensing Resistors (FSRs). These sensors were integrated into two synchronized wireless wearable modules for simultaneous carotid and femoral pulse-wave acquisition and carotid–femoral Pulse Wave Velocity estimation. Hardware synchronization preserved the relative timing between the two modules independently of Bluetooth Low Energy transmission. In 101 volunteers, comparison with SphygmoCor yielded Pearson correlation and Lin concordance coefficients of 0.823, together with a mean paired difference of $0.03\,\mathrm{m\,s}^{-1}$. The mean within-subject standard deviation was $0.63\,\mathrm{m\,s}^{-1}$ for the wearable system and $0.40\,\mathrm{m\,s}^{-1}$ for the reference device. These results support the feasibility of supervised wearable PWV assessment, while indicating that sensor fixation, placement, and waveform-quality control require further improvement.

The main strength of the thesis is the consistent use of application-level information requirements to guide complete biomedical instruments. Despite their prototype maturity and validation limitations, the systems show that information-centred design can reduce complexity and cost while preserving the required measurements, providing a basis for scalable neural interfaces and more accessible cardiovascular monitoring.

Contents

List of Figures **xii**

List of Tables **xx**

Nomenclature **xxiii**

1 Introduction **1**

2 Neural Bioelectronic Systems: Background and State of the Art **3**

2.1 Motivation and Chapter Scope 3

2.2 Neural Activity and Extracellular Spikes 4

2.3 In Vitro Neuronal Cultures 6

2.3.1 Applications 6

2.4 Microelectrode Arrays as the Recording Interface 7

2.5 Requirements for Network-Level Recording 8

2.6 State of the Art 9

2.6.1 Continuous Multichannel Sampling 9

2.6.2 Commercial and Research Platforms 10

2.7 Research Gap and Rationale of the Following Study 11

2.7.1 From Continuous Waveforms to an Event-Driven Representation 11

2.7.2 Objective of Research 13

3 Event-Driven Neural Acquisition Platform 14

3.1 Introduction and Motivation 14

3.2 Design Strategy and System Architecture 16

3.2.1 Event-Driven Neural Acquisition Requirements 16

3.2.2 Event-Driven Acquisition Principle 16

3.2.3 System Overview 18

3.3 Shared Hardware Platform 21

3.3.1 Modular Organization 21

3.3.2 Low-Noise Analog Front End 22

3.3.3 Motherboard 23

3.4 MCU-Based 15-Channel Implementation 25

3.4.1 General Operating Sequence 25

3.4.2 Calibration Phase 27

3.4.3 Event-Driven Acquisition Phase 28

3.4.4 Host Interface 28

3.4.5 Architectural Limitations 29

3.5 FPGA-Based 60-Channel Implementation 30

3.5.1 General Operating Sequence 30

3.5.2 Calibration Phase 30

3.5.3 Parallel Event-Driven Acquisition 32

3.5.4 Host Software and Threshold Estimation 34

3.5.5 Architectural Characteristics and Limitations 36

3.6 Validation Methodology 37

3.6.1 Validation Strategy and Interpretation Boundaries 37

3.6.2 Preliminary Single-Channel AFE Demonstration 38

3.6.3 Common Synthetic-Signal Backend Test Bench 40
3.6.4 MCU Proof-of-Concept Protocol 43
3.6.5 FPGA Validation Campaign 45
3.6.6 FPGA Temporal Alignment and Event Matching 49
3.6.7 Statistical Analysis 51
3.7 Results 52
3.7.1 Preliminary Single-Channel AFE Results 52
3.7.2 MCU Proof-of-Concept Results 54
3.7.3 FPGA-Based Validation Results 56
3.7.4 Output Data Rate 65
3.7.5 Prototype Cost and Accessibility 66
3.8 Discussion 67
3.9 Conclusions 71

4 Pulse-Wave Monitoring: Background 73
4.1 Clinical Motivation and Chapter Scope 73
4.2 From Blood Pressure to the Arterial Pulse Wave 74
4.2.1 Blood Pressure and Pulse-Wave Generation 74
4.2.2 Arterial Wall Mechanics and Wave Propagation 75
4.3 Pulse-Wave Analysis and Pulse-Wave Velocity 76
4.3.1 Information Contained in the Pulse Wave 76
4.3.2 PTT, PAT, and PWV 77
4.3.3 Carotid–Femoral PWV 78
4.4 Non-Invasive Pulse-Wave Acquisition 80
4.4.1 Sensing Modalities and Pressure-Sensor Technologies . . . 80
4.4.2 Instrumentation Requirements 81
4.5 State of the Art in PWV Measurement Systems 82

4.5.1 Clinical Reference and Commercial Systems 83
4.5.2 Portable and Wearable Approaches 84
4.5.3 Current Technological Limitations 84
4.6 Research Gap and Objectives 85

5 Graphene-Based Pressure Sensing Platform for Pulse Wave Acquisition 87
5.1 Introduction and Motivation 87
5.2 Design Strategy and System Overview 88
5.2.1 Pulse Wave Acquisition Requirements 89
5.2.2 System Overview 89
5.3 Sensor Design 91
5.3.1 Material and Working Principle 91
5.3.2 Sensor Structure and Geometry Optimization 91
5.4 Mechanical Integration and Sensor Coupling 93
5.4.1 Pen-Shaped Holder and Pressure-Transfer Interface 93
5.4.2 Sensor Mounting and Electrical Coupling 94
5.4.3 Spring-Assisted Elastic Recovery 95
5.5 Electronic Acquisition System 95
5.5.1 Current-Based Wheatstone Bridge 96
5.5.2 Analog Conditioning Chain 97
5.5.3 Automatic Bridge Balancing 98
5.5.4 Embedded Acquisition Module 99
5.6 Firmware and Host Software 99
5.6.1 Firmware Workflow 100
5.6.2 Bridge-Balancing Routine 100
5.6.3 Host Software and User Interface 100
5.7 Experimental Validation Methodology 101

5.7.1 Sensor-Level Characterization 101
5.7.2 Dynamic Mechanical Characterization 104
5.7.3 Preliminary Pulse-Wave Acquisition 105
5.8 Results 105
5.8.1 Piezoresistive Response and Cyclic Stability 106
5.8.2 Dynamic Mechanical Response 108
5.8.3 Preliminary Pulse-Wave Acquisition 110
5.9 Discussion 112
5.10 Conclusions 114

6 Wearable Device for Pulse Wave Velocity Monitoring 116
6.1 Introduction and Motivation 116
6.2 Design Strategy and System Overview 117
6.2.1 Requirements for Wearable cfPWV Monitoring 118
6.2.2 System Overview 118
6.3 Sensor Selection and Working Principle 120
6.4 Wearable Integration and Sensor Coupling 121
6.5 Electronic Acquisition System 123
6.5.1 Analog Front-End 123
6.5.2 Embedded Unit 124
6.6 System Control and Wireless Communication 126
6.6.1 Operating Sequence and Front-End Calibration 126
6.6.2 BLE Data Transmission and Receiver Interface 127
6.6.3 Host Software and User Interface 128
6.7 Signal Processing and PWV Estimation 129
6.7.1 Signal Pre-processing 129
6.7.2 Fiducial-Point Extraction and PWV Calculation 129

6.8 Validation Methodology 132
6.8.1 Study Design and Population 132
6.8.2 Measurement Procedure and Reference System 132
6.8.3 Statistical Analysis 135
6.9 Results 136
6.9.1 Waveform Quality and PWV Extraction 136
6.9.2 Agreement with the Reference System 136
6.9.3 Repeatability and Age-Stratified Performance 138
6.10 Discussion 138
6.11 Conclusions 140

7 Conclusions **141**
7.1 Main Outcomes of the Thesis 141
7.2 Methodological Considerations 143
7.3 Limitations and Future Perspectives 144

References **146**

Appendix A Supplementary Results **157**
A.1 Neural Acquisition Platform 157
A.1.1 Complete FPGA Study 1 Results 157
A.1.2 Complete FPGA Studies 2 and 3 Results 164
A.1.3 Complete FPGA Study 4 Latency Results 165

List of Figures

2.1 Ionic currents, intracellular action potential, and extracellular spike . 5

3.1 Overview of the event-driven neural acquisition platform. (a) High-level block diagram of the complete system, including the passive MEA, the modular analog front-end, the motherboard implementing signal interfacing and hardware spike detection, the digital control unit, and the host computer for visualization and storage. (b) Standard MCS 60-electrode MEA footprint [1]; (c) Photograph of the modular analog front-end boards. (d) Photograph of the motherboard hosting the MEA interface and event-detection circuitry. (e) STM32 NUCLEO-H723ZG development board [2]; (f) Lattice MachXO2 Breakout Board [3]. 20

3.2 Single-channel diagram of the event-driven acquisition platform: (a) simplified AFE diagram, showing the multiple amplifying and filtering stages for signal readout; (b) motherboard event-detection diagram, consisting of two polarity paths for detecting both positive and negative spikes; and (c) non-inverting Schmitt trigger topology used for hardware threshold detection. 23

3.3 Operation and data formatting of the MCU-based 15-channel implementation. (a) Simplified workflow of initialization, calibration, acquisition, and host interaction. (b) Binary activity word representing one 1 ms time bin. (c) Packet structure used to transmit 500 ms blocks of event activity to the host computer. 26

3.4 Operating sequence of the FPGA-based implementation, showing the distribution of initialization, optional calibration, parallel event acquisition, buffering, communication, and host-side processing across the motherboard, FPGA, and MATLAB host. 31

3.5 Parallel event conditioning in the FPGA backend. (a) Each asynchronous motherboard event line $e_i(t)$ is sampled through a five-stage register. A three-out-of-five majority rule produces the filtered level $s_i[k]$, while a one-cycle delay provides $s_i[k-1]$. Their combination, $p_i[k] = s_i[k] \wedge \neg s_i[k-1]$, generates one pulse on each accepted rising transition. This pulse is OR-accumulated into the sticky occupancy bit $m_i[w]$, which is cleared after the current acquisition window has been captured. The datapath is replicated independently for all 60 channels. (b) An isolated transient produces only one high sample and is rejected, whereas a sustained event raises the five-sample population count $c_i[k]$ to three, asserts $s_i[k]$, and generates a single-cycle rising-edge pulse. 32

3.6 Synthetic-signal backend validation setup. (a) Single spike waveform generated in MATLAB and reproduced by the NI USB-6259. (b) Complete FPGA-based experimental setup, including the DAQ, FPGA control unit, and motherboard. (c) Signal-distribution PCB used to replicate the DAQ output across the 60-channel motherboard interface with the standard MEA footprint. (d) Passive bridge modules mounted in place of the analog front ends, allowing the synthetic post-AFE signal to be injected directly into the motherboard detection path. 41

3.7 Temporal alignment using a Barker-13 preamble sequence, with each spike representing a 1, while noise represents a 0. 50

3.8 Preliminary characterization of the single-channel analog front end. (a) Zero-input electrical recording, showing the background noise of the electronic acquisition chain. (b) Response to controlled microvolt-scale artificial pulse train. (c) Representative *in vitro* recordings obtained from four selected electrodes of a passive MEA. 53

3.9 Results of the MCU-based 15-channel proof-of-concept validation. (a) Offline overlay of a representative 400 ms synthetic input segment at $SNR_{lin} = 10$ and $f_r = 10\,Hz$, showing the noisy input waveform in black, the known spike occurrences as blue circles, and the reconstructed events as red arrows. The overlay was generated during post-processing for qualitative visualization and does not constitute an absolute temporal-alignment measurement. (b) Representative host-side raster plot obtained during a 10 s acquisition with all 15 channels stimulated concurrently. (c) Channel-averaged F_1 score as a function of the dimensionless linear amplitude ratio SNR_{lin}, evaluated at $f_r = 10\,Hz$. (d) Channel-averaged F_1 score as a function of firing rate for selected values of SNR_{lin}. For representation clarity, not all the acquisition sessions are shown. 55

3.10 Detection operating envelope obtained in FPGA Study 1. The heatmaps report the grand-average F_1 score across the tested $SNR_{lin} - f_r$ space for each threshold multiplier K. The white dashed contour identifies the adopted $F_1 \geq 0.75$ operating threshold. Asterisks mark conditions in which the corresponding K value produced the statistically highest session-level F_1 score after FDR correction. . . 58

3.11 Inter-session repeatability in FPGA Study 1, expressed as the CV of the session-level F_1 score. The white dashed contour marks the predefined $CV = 10\%$ threshold. 59

3.12 Mean temporal jitter in FPGA Study 1 as a function of SNR_{lin} and threshold multiplier K. The analysis only included firing rates up to 100 Hz and to conditions with effective detection. Error bars indicate the standard deviation. 60

3.13 Detection performance in FPGA Study 2 as a function of spike amplitude after a single calibration at $A_{ref} = 150\,mV$. Each curve reports the channel-averaged F_1 score for one experimental session. The calibration-derived noise estimate and DAC thresholds were held fixed throughout the amplitude sweep. 62

3.14 Detection performance in FPGA Study 3 as a function of mean firing rate under Poisson-distributed spike timing. Colored curves report the channel-averaged F_1 score for the individual sessions, while the dashed black curve indicates the session average. The progressive reduction at high mean firing rates is primarily associated with closely spaced events affected by the occupancy-window and refractory constraints. 63

3.15 Event-driven system latency in FPGA Study 4 as a function of $\mathrm{SNR}_{\mathrm{lin}}$ and threshold multiplier K. Bars report the mean of the ten session-level mean latencies, and error bars indicate their standard deviation across sessions. Brackets compare $K = 3$ and $K = 4$ within each SNR condition using exact Wilcoxon signed-rank tests after FDR correction; an asterisk denotes a significant difference and n.s. a non-significant difference. The effect of SNR was evaluated separately within each K condition and is reported in the text and in Appendix A.1.3. The $\mathrm{SNR}_{\mathrm{lin}} = 20$ measurements involved a different comparator-output transition and therefore require separate physical interpretation. 64

4.1 Schematic of arterial pulse-wave generation and propagation. Left-ventricular ejection generates a pressure wave that travels along the arterial tree and is measured non-invasively at the carotid and femoral sites. Pulse transit time (PTT) is the delay between the feet of the two waveforms, and carotid–femoral pulse wave velocity (PWV) is estimated from the corrected path length and PTT. 75

4.2 Arterial pulse-waveform phases and landmarks 77

4.3 Comparison of sequential ECG-referenced and simultaneous one-step carotid–femoral pulse wave velocity acquisition. 79

5.1 The pressure sensor is interfaced to the circular PCB, which performs signal readout; the external embedded module, referred to as the main board, performs further signal processing, digitization, and communication; and the host software running on a computer provides real-time visualization and data storage through its GUI. . . 90

5.2 Working principle and geometry of the custom rGO-R1 flexible-PCB sensors. (a) Conceptual piezoresistive mechanism, in which applied pressure increases conductive contact and reduces resistance. (b) Rectangular sensor section with labeled polyimide, metal, and rGO-R1 layers; current path is shown with a blue line; current inlet and outlet are indicated with red and green arrows. (c) Cross-shaped sensor section using the same color code; terminals A–D are indicated and the illustrated current path extends from B to C. (d,e) von Mises stress distributions for the rectangular and cross-shaped geometries under a central 0.5 N load applied through a 2 mm rounded tip. . . . 92

5.3 Mechanical integration of the cross-shaped sensor. (a) Photograph of the assembled pen-shaped probe. (b) Cross-sectional view showing the skin-contact tip, pressure-transfer shaft, lateral supports, fastening disks, elastic-return spring, spacer, circular analog-front-end PCB, and spring-loaded electrical connectors. (c) Three-dimensional rendering of the detachable sensor-to-PCB interface. 94

5.4 Electronic acquisition architecture. (a) Simplified block diagram of the overall acquisition chain. The probe-side circuitry converts the sensor resistance variation into a differential voltage, the main board provides filtering and programmable gain, and the MCU controls bridge balancing, digitizes the two channels, and communicates with the host computer. (b) Detailed schematic of the current-based Wheatstone bridge, readout, and balancing circuit. 96

5.5 Bridge balancing algorithm (a) and host-software UI (b) for the rGO-R1 platform. The balancing routine iteratively updates $V_{O,A}$ and $V_{O,B}$ from the measured residual currents, while the UI provides connection, calibration, gain selection, two-channel visualization, acquisition control, and data storage. 101

5.6 Experimental validation at sensor, system, and physiological levels. (a) SEM images used to classify the rGO-R1 coating. (b) Shaker–vibrometer setup used to compare the electrical output with controlled diaphragm displacement. (c) Carotid pulse-wave acquisition setup. . 102

5.7 Sensor-level response of the rectangular flexible-PCB devices. The plots report relative resistance variation versus pressure (a), sensitivity in the low-pressure range (b), and the response during 1000 cycles at 0.5 kPa (c). Sensitivity error bars denote the standard error of the mean across at least five samples. The solid lines in (b) represent linear fits used to estimate sensitivity in the low-pressure regime. . . 106

5.8 Dynamic response of the rectangular and cross-shaped sensors. For both rectangular (a) and cross-shaped (b) geometries, the left plot reports output voltage versus measured displacement, while the right plots compare normalized sensor output and reference displacement at representative amplitudes below and at the response plateau. The output voltage has opposite polarity with respect to the imposed displacement because increasing deformation decreases the rGO-R1 sensor resistance through the implemented readout chain. 109

5.9 Representative physiological recordings. (a) One-minute femoral acquisition. (b) Five-second simultaneous carotid and femoral acquisition using two probes connected to the same main board. (c) Qualitative comparison of pulse-wave morphology between the rGO-R1 sensor and SphygmoCor applanation tonometry, using separately acquired representative waveforms. 111

6.1 System architecture of the wearable PWV platform. Two body-worn analog front-end (AFE) modules, positioned over the carotid and femoral arteries, integrate an FSR, resistance-to-voltage conversion, analog conditioning, 12-bit digitization, and BLE transmission. Before placement, the modules are temporarily coupled through a hardware interface to establish a common timing reference. During acquisition, the synchronized data streams are transmitted to a dedicated receiver and forwarded to a host computer, where a graphical user interface (GUI) supports real-time visualization, data storage, and subsequent PTT and PWV estimation. Adapted from [4]. 119

6.2 Wearable integration of the PWV monitoring system [4]. (a) Render of the external side of the enclosure, including the region for the hardware synchronization interface and openings for ON/OFF switch and USB cable. (b) Render of the sensor-side view, showing the FSR seat at the skin-facing region, opening for the sensor to reach the readout electronics, and lateral flanges for the elastic band. (c) Photograph of the assembled wearable modules. (d) Example of carotid and femoral module positioning on a participant during the clinical validation study. (e) Coupled enclosures for module synchronization. 122

6.3 Simplified analog front-end of each wearable sensing module. The FSR and the SPI-controlled digital potentiometer form one branch of a 3 V Wheatstone bridge, while two fixed resistors form the opposite branch. The differential bridge output is processed by an instrumentation amplifier referenced to 1.5 V. Third-order low pass and high-pass stages define the approximate 0.2 Hz to 20 Hz analog passband before ADC digitization at 1 kHz. 124

6.4 Operating sequence of each wearable sensing node. Module decoupling generates the hardware event used to initialize the local acquisition time base. The host can then trigger front-end calibration or timer-driven sampling. Samples are grouped into eight-sample packets and transmitted to the receiver through BLE. 126

6.5 Graphical user interface screenshot. 128

6.6 Processing steps used to identify the ITP: (a) extraction of reference minima from the 1.5 Hz low-pass-filtered waveform; (b) identification of the local pulse-wave foot; (c) calculation of the tangent at the point of maximum derivative; and (d) identification of the ITP from the intersection between the tangent and the horizontal line passing through the foot level. Numerical time and amplitude scales are omitted because they vary among acquisitions. Adapted from [4]. . 131

6.7 Alternating protocol used for the comparison between the proposed wearable system and SphygmoCor. Measurements with the two devices were alternated to minimize the temporal interval between the methods, while operator assignments were alternated across consecutive participants to reduce operator-dependent bias. Adapted from [4]. 134

6.8 Agreement between the proposed wearable system and SphygmoCor: (a) correlation plot of the PWV values obtained with the two devices; and (b) Bland–Altman plot showing the paired differences, mean bias, and limits of agreement. Adapted from [4]. 137

A.1 Alternative representations of the FPGA Study 1 detection results. (a) Dependence of the F_1 score on $\mathrm{SNR}_{\mathrm{lin}}$, highlighting the progressive shift of the detection boundary as K increases. (b) Dependence on firing rate, showing the high-rate degradation under conditions in which the calibration interval was strongly contaminated by the regular spike train. 158

List of Tables

2.1 Extracellular neural-signal characteristics 6

3.1 Principal design requirements and corresponding architectural choices of the event-driven neural acquisition platform. 17

3.2 Main specifications of the event-driven neural acquisition platform. The MCU and FPGA entries describe successive control-unit implementations. 18

3.3 Experimental design of the FPGA validation campaign. Conditions in Studies 1–3 were repeated in five independent sessions. 46

3.4 Representative MCU proof-of-concept results at $SNR_{lin} = 10$, aggregated over firing rates up to 300 Hz. 54

3.5 Summary of the FPGA-based validation results. 57

3.6 Comparison between the event-driven output rates and ideal continuous 12-bit acquisition. The continuous rates assume ideal sample packing and exclude protocol overhead. The $20\,kS\,s^{-1}$ and $50\,kS\,s^{-1}$ cases are analytical extrapolations. 66

4.1 Representative non-invasive sensing and timing approaches for pulse-wave and PWV assessment. 81

4.2 Qualitative comparison of pressure-sensor technologies for non-invasive pulse-wave acquisition. 82

4.3 Representative systems for PWV and arterial-stiffness assessment. . 83

5.1 Main specifications of the rGO-R1 pulse-wave acquisition platform [5]. 90

5.2 Static sensitivity values calculated from the low-pressure response of the flexible PCB pressure sensors. 107

5.3 Relative resistance response at 0.5 kPa before and after 1000 press-release cycles. 108

5.4 Representative dynamic mechanical characterization results for rectangular and cross-shaped sensors. 110

6.1 Main specifications of the wearable PWV monitoring system [4]. . . 120

6.2 Main characteristics of the validation cohort. Continuous variables are reported as mean ± standard deviation. 133

6.3 Age-stratified difference between the proposed wearable system and SphygmoCor. 138

A.1 Complete FPGA Study 1 detection results for $K = 3$. 160

A.2 Complete FPGA Study 1 detection results for $K = 4$. 161

A.3 Complete FPGA Study 1 detection results for $K = 5$. 162

A.4 FPGA Study 1 conditions exceeding the $CV = 10\%$ repeatability threshold. 162

A.5 FPGA Study 1 conditions showing a statistically significant effect of K on F_1. Reported p-values are false-discovery-rate corrected. . . . 163

A.6 Mean temporal jitter in FPGA Study 1 as a function of $\mathrm{SNR}_{\mathrm{lin}}$ and K. Values are mean ± standard deviation in milliseconds. A dash indicates that effective detection was not achieved. 164

A.7 Complete detection and temporal results for FPGA Study 2. Calibration was performed at $A_{\mathrm{ref}} = 150\,\mathrm{mV}$, and the resulting thresholds were held fixed throughout the amplitude sweep. 164

A.8 Complete detection and temporal results for FPGA Study 3 under Poisson-distributed spike timing. 165

A.9 Per-session FPGA-internal latency results. The final row reports the mean and standard deviation of the 15 session means and the RMS of the 15 within-session standard deviations. 167

A.10 Per-session event-driven system latency for $K = 3$. The summary row for each SNR reports $\overline{\mu}_{\text{sess}} \pm s_{\mu}$ in the mean column and σ_{RMS} in the within-session SD column. 168

A.11 Per-session event-driven system latency for $K = 4$. The summary row for each SNR reports $\overline{\mu}_{\text{sess}} \pm s_{\mu}$ in the mean column and σ_{RMS} in the within-session SD column. 169

A.12 Friedman tests evaluating the effect of SNR on the session-level mean event-driven system latency. The two tests, corresponding to $K = 3$ and $K = 4$, formed one predefined family ($m = 2$) for Benjamini–Hochberg false-discovery-rate correction. 169

A.13 Exact paired Wilcoxon signed-rank comparisons among SNR conditions using the session-level mean event-driven system latency. The six comparisons, comprising three SNR pairs for each value of K, formed one predefined family ($m = 6$) for Benjamini–Hochberg false-discovery-rate correction. 170

A.14 Exact paired Wilcoxon signed-rank comparisons between $K = 3$ and $K = 4$ using the session-level mean event-driven system latency. The three comparisons, corresponding to the three SNR conditions, formed one predefined family ($m = 3$) for Benjamini–Hochberg false-discovery-rate correction. 170

Nomenclature

Acronyms / Abbreviations

ADC Analog-to-digital converter

AFE Analog front end

BLE Bluetooth Low Energy

BP Blood pressure

cfPWV Carotid–femoral pulse wave velocity

CMOS Complementary metal–oxide–semiconductor

CV Coefficient of variation

DAC Digital-to-analog converter

DAQ Data-acquisition device

DIV Days in vitro

ECG Electrocardiogram

EIS Electrical impedance spectroscopy

EIT Electrical impedance tomography

FDR False discovery rate

FFC Flexible flat cable

FN False negative

FP False positive

FPGA Field-programmable gate array

FSR Force-sensing resistor

GPIO General-purpose input/output

GUI Graphical user interface

INA Instrumentation amplifier

ISI Inter-spike interval

ITO Indium tin oxide

ITP Intersecting tangent point

MAD Median absolute deviation

MCU Microcontroller unit

MIT Magnetic induction tomography

PAT Pulse arrival time

PCB Printed circuit board

PET Polyethylene terephthalate

PPG Photoplethysmography

PTT Pulse transit time

PWA Pulse wave analysis

rGO-R1 Reduced graphene oxide functionalized with triethylene glycol amine (R1 formulation)

RMS Root mean square

SEM Scanning electron microscopy

SNR Signal-to-noise ratio

SoC System on chip

SPI Serial Peripheral Interface

TIA Transimpedance amplifier

TP True positive

UART Universal asynchronous receiver-transmitter

USB Universal Serial Bus

Chapter 1

Introduction

The core motivation of this thesis is the possibility of translating advances in sensing, electronics, and embedded processing into healthcare technologies that are more accessible, efficient, and responsive to real application needs. Achieving this objective does not depend only on improving an individual sensor or processing algorithm. It requires the complete instrument to be built around the biological information that must ultimately be obtained.

Biological information is not directly available to an electronic system. It must be accessed through a physical interface, converted into an electrical signal, conditioned within a suitable amplitude and bandwidth range, acquired with adequate temporal fidelity, and transformed into quantities that are meaningful for the intended application. The quality and impact of a biomedical device therefore depend on the complete acquisition chain rather than isolated sensor, amplifier, or algorithm performance.

This thesis uses this system-level and application-driven perspective for two different biomedical application domains: multichannel extracellular recording from *in vitro* neuronal cultures and non-invasive arterial pulse-wave monitoring. The two domains present numerous differences in their biological processes, sensing interfaces, signal characteristics, and experimental settings. They also impose different information requirements. For some applications, neural network monitoring may be based primarily on the occurrence, timing, and spatial distribution of extracellular spikes, whereas pulse wave assessment requires preservation of the pulse-wave morphology.

The purpose of the thesis is not to combine these application domains at the physiological level or to propose a universal biomedical instrument. Rather, it is to investigate how application-level information requirements can guide the design, integration, and validation of biological acquisition systems, with the broader aim of contributing to healthcare technologies that are more accessible and efficient.

The thesis addresses two research questions:

1. Can neural acquisition instrumentation be made more accessible and efficient when network-level activity, rather than extracellular waveform morphology, is the principal information of interest?

   This question is addressed in Chapters 2 and 3 through the development and staged validation of a modular platform that performs hardware-level spike extraction and represents neural activity through channel identity and event timing. This represents the principal research line of this thesis.

2. Which sensing and system-level solutions can support accessible, usable, and robust non-invasive arterial pulse-wave monitoring?

   This question is addressed in Chapters 4, 5, and 6 through the development of a custom piezoresistive sensing platform and a subsequent wearable system for carotid–femoral pulse-wave acquisition and PWV assessment.

The structure of this thesis is organized as follows.

Chapter 2 introduces extracellular neural activity, *in vitro* cultures, MEAs, conventional acquisition architectures, and the application-specific gap motivating event-driven readout. Chapter 3 presents the neural acquisition platform, from the 15-channel implementation reported in [6] to the subsequent 60-channel architecture, together with the staged validation of design feasibility and analytical performance.

Chapter 4 introduces arterial pulse-wave monitoring and the requirements of carotid–femoral pulse wave velocity (cfPWV) assessment. Chapter 5 presents the custom graphene-based pressure-sensing platform reported in [5], from its development and validation from sensor-level response to preliminary physiological acquisition. Chapter 6 presents the subsequent wearable system reported in [4] and its comparison with a clinical reference device.

Finally, Chapter 7 synthesizes the principal outcomes, methodological implications, limitations, and future development directions.

Chapter 2

Neural Bioelectronic Systems: Background and State of the Art

2.1 Motivation and Chapter Scope

Neural acquisition systems are devices that provide an electrical interface through which neural activity can be recorded, analysed, and, in bidirectional configurations, used to guide stimulation. *In vitro* neuronal cultures, grown on microelectrode arrays (MEAs), offer an experimentally accessible model in which network activity can be observed over multiple sites for extended periods. Neural cultures support studies of network development and synchronization, pharmacological and neurotoxicity assays, disease-oriented cellular models, and closed-loop experiments in which recorded activity affects the subsequent stimulation or environment [7–10].

The electrical activity of neurons is commonly accessed through a MEA. Electrodes distributed beneath the culture sense local extracellular potential variations relative to a reference electrode. These signals are amplified, filtered, and acquired simultaneously across multiple channels, allowing the temporal and spatial organization of network activity to be reconstructed.

The resulting instrumentation problem is defined by a mismatch of scales. At the electrode level, the available signals are brief, typically microvolt-level extracellular spikes superimposed on electrode noise, baseline drift, and interference. The biological questions, however, are often formulated at the network level: which

recording sites are active, how frequently they fire, whether their activity is organized into bursts, and how events propagate or synchronize across the culture. The acquisition system must therefore preserve the information required by the experiment, taking into account the characteristics of the recorded signals.

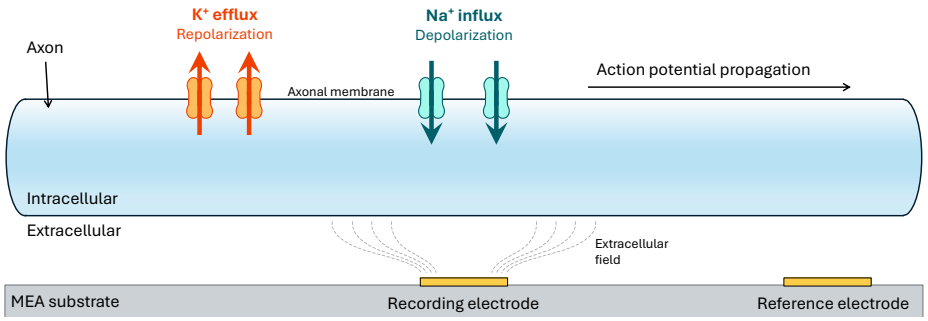
This chapter introduces neural electrical activity and the conventional methods used for its acquisition. Finally, the application-specific gap is reported, motivating the acquisition strategy developed in Chapter 3.

2.2 Neural Activity and Extracellular Spikes

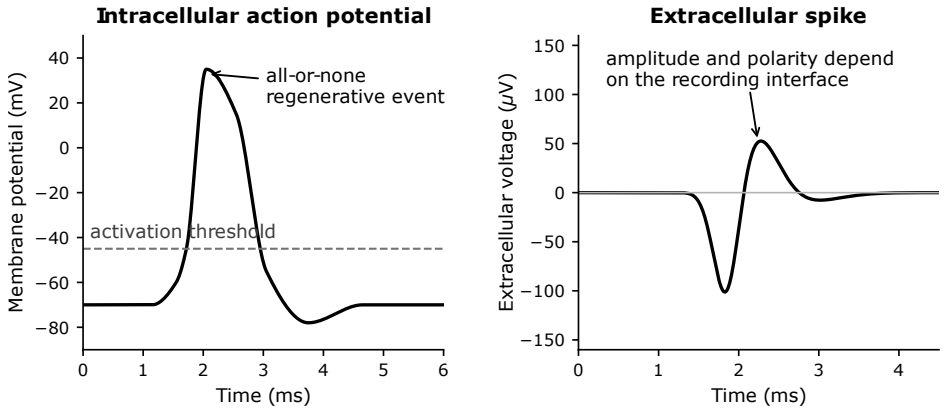
Neurons communicate through changes in membrane potential. When a sufficiently strong depolarization reaches the activation threshold, voltage-dependent ionic currents produce a regenerative action potential. The associated ionic currents are illustrated in Figure 2.1(a). Once initiated, this response is approximately all-or-none: increasing a supra-threshold stimulus does not produce a proportionally larger action-potential amplitude [11]. Information is instead carried primarily through the occurrence and timing of action potentials, their firing rate and temporal pattern, and their distribution across a population of neurons.

An extracellular electrode measures a local potential relative to a reference, not an attenuated copy of the intracellular membrane voltage. The observed waveform results from the spatial and temporal distribution of transmembrane currents and is shaped by neuronal geometry, electrode position, cell–electrode coupling, electrode impedance, and the transfer function of the recording chain [8, 12]. Extracellular spikes may consequently be monophasic, biphasic, or triphasic and may present a dominant positive or negative component.

The distinction between the intracellular action potential and the extracellular spike is summarized in Figure 2.1. The intracellular signal is a regenerative membrane-potential excursion with an amplitude of several tens of millivolts. By contrast, an extracellular electrode records a much smaller and shorter voltage transient, typically in the microvolt range, derived from the occurrence of action potentials. Its amplitude, polarity, and morphology are not fixed, but depend on the position of the active cells, their orientation relative to the electrode, the electrode–electrolyte interface, and the transfer characteristics of the recording chain.



(a) Ionic currents and extracellular recording.



(b) Intracellular action potential and extracellular spike.

Fig. 2.1 Conceptual illustration of neural electrical activity and extracellular recording. (a) Ionic currents associated with an action potential propagating along an axon. Sodium influx contributes to depolarization, whereas potassium efflux contributes to repolarization in the region behind the advancing action potential. A recording microelectrode on the MEA substrate senses the local extracellular potential relative to a reference electrode. The dashed connection schematically indicates extracellular potential sensing and does not represent a physical field line or a current path. (b) Comparison between an intracellular action potential and an extracellular spike. The intracellular event has a millivolt-scale all-or-none waveform, whereas the extracellular spike is a shorter, microvolt-scale signal whose amplitude, polarity, and morphology depend on the recording interface. Panel (a) is schematic and not to scale; waveform scales in (b) are representative.

The principal extracellular signal characteristics relevant to the present work are also reported in Table 2.1 [8, 13–15].

The acquisition chain must therefore resolve small and brief signals without assuming a single amplitude, waveform, or baseline condition for the complete array. At the same time, it must preserve the relative timing of activity across channels, because network-level information is carried not only by isolated spikes but also by their temporal and spatial organization.

Table 2.1 Representative characteristics of extracellular spikes in *in vitro* MEA recordings and related instrumentation requirements.

Characteristic	**Representative range**	**Instrumentation requirement**
Amplitude	Tens–hundreds of µV	Low noise, high input impedance, and adequate gain.
Morphology	Mono-, bi-, or triphasic; either polarity	Bipolar signal handling and spike-band preservation.
Duration	Approximately 0.5 ms to 2 ms	Adequate bandwidth and temporal resolution.
Spike band	Typically 300 Hz to 3000 Hz	Band-pass conditioning suitable for spike detection.
Firing rate	Sub-hertz to tens of hertz; higher during bursts	Stable acquisition under sparse and dense activity.
Spatial activity	Local, propagating, or synchronous	Synchronized multichannel acquisition.
Channel variability	Different offsets, noise levels, and amplitudes	Per-channel characterization and configuration.
Noise and artifacts	Drift, interference, coupling	Filtering, shielding, sufficient dynamic range.

2.3 In Vitro Neuronal Cultures

Dissociated neuronal cultures are obtained by isolating neural cells, seeding them on a suitable substrate, and maintaining them under controlled conditions. During maturation, neurons form synaptic connections, and develop spontaneous electrical activity. The resulting network dynamics include isolated spikes, repeated bursts, changes in firing rate, propagation between regions, and varying degrees of synchrony [15, 8]. Cultures do not reproduce the complete anatomical and physiological context of an intact organism, but provide direct access to the preparation, compatibility with microscopy and pharmacological manipulation, and repeatable observation over long periods.

2.3.1 Applications

In-vitro MEA recordings are commonly used to examine network maturation and plasticity, quantify the response to drugs and chemicals, characterize disease-specific

or patient-derived cellular models, and assess the effects of electrical or optical stimulation [9, 8]. Information required for these applications often includes the number of active electrodes, mean firing rate, burst rate and duration, inter-burst interval, synchrony, recruitment, and spatial propagation. These quantities reduce the recorded activity to a spatiotemporal description of how the network evolves and responds to an experimental intervention. Other studies require instead the complete waveform, for example to perform morphology-based spike sorting, distinguish multiple units recorded by one electrode, analyse low-frequency field potentials, inspect artifacts, or repeat filtering and detection retrospectively [14, 16].

Closed-loop experiments extend recording from observation to interaction. The system acquires neural activity, estimates a relevant network state, and uses that estimate to select a subsequent stimulation or environmental action. Such approaches have been used to regulate bursting in cortical cultures and to study activity-dependent adaptation under controlled feedback [10, 17]. They place additional emphasis on predictable latency, synchronized multichannel acquisition, reliable detection of concurrent spikes, and stable operation over extended sessions.

A related emerging direction couples living neuronal cultures to simulated or physical environments. In these bio-hybrid systems, neural activity may be mapped to an external action, while sensory or task-related information is returned through stimulation. Recent demonstrations have explored whether such interfaces can support adaptive network behaviour in simplified tasks [18, 19]. These studies remain methodologically and conceptually distinct from established pharmacological or network-analysis assays, but they reinforce the need for instrumentation that can expose multichannel activity in real time without introducing poorly controlled timing or communication bottlenecks.

2.4 Microelectrode Arrays as the Recording Interface

An *in vitro* MEA consists of multiple microelectrodes exposed through an insulating substrate and connected to external pads or integrated readout circuits. Neurons located near an electrode generate extracellular potential variations relative to a reference. Each site senses a local mixture of nearby current sources; an MEA therefore samples the spatial organization of network activity rather than establishing a one-to-one correspondence between an electrode and a neuron [7, 13].

Passive MEAs are only constituted by electrodes and conductive tracks, leaving amplification, filtering, digitization, and stimulation to external instrumentation. Their main advantages are experimental accessibility, optical compatibility, interchangeable biological preparations, and the possibility of modifying the acquisition electronics independently of the sensor substrate.

On the other hand, active complementary metal–oxide–semiconductor (CMOS) MEAs integrate switches, amplifiers, and readout circuits close to a dense electrode matrix. They can provide thousands or tens of thousands of sites, flexible electrode selection, and high spatial resolution [20, 21]. This integration reduces parasitic effects and enables measurements that are not practical with conventional passive arrays, but it also creates a specialized design space involving custom fabrication, packaging, routing, and large-scale data handling.

The electrode–electrolyte interface is frequency dependent and contributes impedance and noise to the recording path. Together with parasitic capacitance and amplifier input impedance, it affects signal attenuation and bandwidth. Electrode half-cell potentials, slow drift, and environmental interference can be substantially larger than the extracellular spikes of interest. Variations in electrode condition, cell adhesion, and local network geometry also produce different baselines, noise levels, and spike amplitudes across the array [8].

Mechanical and experimental factors are equally relevant. The interface must provide reproducible electrical contact, preserve access to the biological preparation, and allow stable operation without introducing excessive coupling from nearby digital electronics or power supplies.

2.5 Requirements for Network-Level Recording

The preceding considerations define the following main acquisition requirements:

- Analog interfacing and signal conditioning. The analog front end (AFE) must amplify microvolt-scale extracellular signals while minimizing electrode loading and avoiding saturation. Its input impedance should substantially exceed the electrode–electrolyte impedance throughout the recording band, considering both input resistance and capacitance. The analog passband must attenuate electrode offsets, slow drift, and out-of-band noise while retaining

spike-related components [22]. Noise should be specified as an input-referred RMS voltage together with the measurement bandwidth and input conditions. As commercial references within the MEA2100 family, the MEA2100-System specifies an input impedance of 1 GΩ in parallel with 10 pF and a default passband of 1 Hz–3 kHz [23]. The MEA2100-Mini specifies typical input noise of 0.7 μV RMS over 1 Hz–3.5 kHz, with inputs connected to ground [24].

- Timing and latency. Propagation and synchrony analyses depend on relative delays between recording sites. All channels should therefore share a common temporal reference, channel-dependent acquisition skew should be minimized, and system behaviour should remain bounded under concurrent activity. Some application, such as closed-loop protocols, require low and predictable latency to estimate the network state.
- Operational stability. The system must support the duration and operating conditions of the experiment. Long-term assays require stable acquisition, reliable communication, manageable storage, and reproducible electrical contact. The physical interface should also preserve access to the culture and limit interference from digital circuits and power supplies.
- Information representation. The information representation should match the experimental objective. Preserving the complete voltage waveform maximizes post-processing flexibility but imposes the highest conversion, communication, and storage requirements. A more compact representation may be appropriate when the analysis depends primarily on the occurrence and distribution of spikes.

2.6 State of the Art

2.6.1 Continuous Multichannel Sampling

The conventional acquisition system architecture relies on continuous signal sampling. This is the appropriate general-purpose solution when the waveform itself is part of the measurement. It supports retrospective filtering, threshold adjustment, spike sorting, artifact inspection, low-frequency analysis, and the application of processing

methods that were not specified when the experiment was performed. Its principal cost is that data transfer, storage, and processing scale directly with the number of channels, the selected bandwidth, and the recording duration. For N_{ch} channels, sampling frequency f_s, and b transmitted bits per sample, the ideal raw data rate is

$$R_{\mathrm{cont}} = N_{\mathrm{ch}} f_s b. \tag{2.1}$$

For example, 60 channels sampled at $10\,\mathrm{kS\,s^{-1}}$ with 12-bit resolution produce an ideal payload of $7.2\,\mathrm{Mbit\,s^{-1}}$, before sample alignment, framing, metadata, or communication overhead. This corresponds to an ideal raw data volume of 3.24 GB per hour, or approximately 1.09 TB over 14 days of uninterrupted recording.

Although storage capacity is relatively inexpensive, retaining continuous recordings also entails data transfer, processing, backup, and maintenance requirements, whether local infrastructure or cloud services are used. These demands become more relevant as channel count, recording duration, and the number of experiments increase.

2.6.2 Commercial and Research Platforms

Commercial MEA systems provide mature combinations of amplification, digitization, stimulation, synchronization, and analysis software. Passive-MEA platforms such as the MEA2100 family preserve multichannel voltage waveforms and support a broad range of recording and stimulation protocols [23]. Multiwell systems such as the Maestro family extend this model toward standardized and higher-throughput pharmacological and toxicological assays [25].

Accessibility is also influenced by instrumentation cost. Published studies have reported complete commercial 60-channel MEA systems in the USD 30,000–50,000 range, with representative acquisition configurations exceeding USD 35,000; later work similarly described commercial *in vitro* MEA platforms as costing several tens of thousands of euros [26–28]. These historical and configuration-dependent estimates should not be interpreted as current quotations, but they give an idea of the economic burden caused by general-purpose instrumentation, which may not be accessible to all prospective users.

Other commercial and research platforms address different priorities. Integrated front ends, such as the RHD series, reduce the effort required to implement a low-noise multichannel recorder while retaining continuous digitized waveforms [29]. Open or custom systems emphasize accessibility, modification, and real-time control; representative examples demonstrate that multichannel recording and stimulation can be implemented with modular hardware and user-defined software [26, 17, 28]. High-density CMOS MEAs instead move amplification, routing, and electrode selection onto the sensor chip to reach much larger electrode counts and finer spatial sampling [20, 21].

Despite their differences, these systems share a defensible design principle: the acquisition stage preserves the waveform so that the same instrument can support many analyses. This is essential when the relevant features are unknown in advance or when voltage morphology is itself an experimental endpoint. It also means, however, that the complete sampled waveform is acquired even when subsequent analysis reduces it to a sequence of detected spikes.

2.7 Research Gap and Rationale of the Following Study

2.7.1 From Continuous Waveforms to an Event-Driven Representation

In many network-level experiments, continuous extracellular voltage traces are not the final information product. The recorded signals are filtered, spikes are detected, and subsequent analyses operate on a reduced spike representation containing the occurrence times and spatial origin, often detecting events simply with a threshold. Firing rate, burst timing, neuronal recruitment, propagation patterns, and synchrony can then be derived from the resulting spatiotemporal activity map [8, 16]. When these quantities constitute the principal observables, continuous acquisition preserves waveform information that may not be used in the final analysis.

This consideration motivates the change of paradigm adopted in this work. Instead of digitizing and transmitting every sample, the analog signal from each channel can be conditioned and processed to detect candidate spikes at a hardware level using a

simple threshold-based method. Each accepted event can then be simply represented by its electrode ID and occurrence time.

Threshold-based spike extraction is widely used both in offline processing and in real-time multichannel detection systems because of its comparatively low computational complexity and straightforward implementation [16, 30, 31]. Channel-specific thresholds are required because the measured extracellular waveform may differ among channels, depending on factors such as neuron–electrode distance, electrode impedance, tissue coupling, and local noise conditions [8, 12].

Event-driven extraction removes continuous amplitude samples and replaces them with a representation of detected activity. The communication and storage rates may scale with the number of detected events rather than with the sampling frequency and channel count. As an alternative, detected events can be accumulated within fixed-duration time windows and transmitted as channel-occupancy records. The resulting time–channel representation is directly compatible with raster visualization, firing-rate estimation, burst analysis, and low-latency closed-loop processing. These properties are also particularly relevant for long recordings, high electrode counts, and low-latency experiments [10, 32, 33].

This reduction entails important drawbacks. Detection performance depends on analog conditioning, noise estimation, threshold calibration, and the stability of the electrode–tissue interface. More importantly, in a strict event-only representation, discarded waveform morphology cannot be reconstructed. Furthermore, given the relatively low cost of data storage, the savings offered by an event-based representation may provide only a limited advantage over continuous waveform recording. Event-driven acquisition should consequently not be regarded as a universal replacement for continuous recording, but as an application-specific alternative for experiments in which spatiotemporal information is the principal feature of interest.

Integrated research systems have demonstrated activity-dependent neural readout using level-crossing and spike-triggered architectures [34, 35]. More recently, the GAIA device extended this concept to a monolithic 64×64 event-based microelectrode array that generates asynchronous address events at pixel level [36]. These implementations demonstrate the technical feasibility of event-based neural sensing, but they are primarily conceived as custom integrated circuits, implantable interfaces, or monolithic active arrays.

2.7.2 Objective of Research

The objective of this research is the development of an accessible, ultra-low data-rate, efficient MEA acquisition platform capable of performing channel-adaptive event extraction directly in hardware for applications targeting spatiotemporal information as their main functional feature.

The literature reviewed above establishes the feasibility of event-based neural readout through custom integrated architectures [34–36]. Building on these approaches, this work investigates a modular implementation that represents detected activity through channel identity and temporal information.

The study addresses the following connected research questions:

1. Can a custom analog front end acquire extracellular activity from neuronal cultures on a standard passive MEA with gain and noise performance suitable for spike detection?

2. Can hardware threshold crossings be encoded in a compact multichannel representation that preserves channel identity and event timing at a specified temporal resolution? Its operating limits require explicit characterization, since threshold-based spike detection depends on signal quality, noise estimation, and threshold selection [16, 31].

3. To what extent do the retained spatiotemporal information and measured system performance meet the requirements of applications primarily concerned with the timing and spatial distribution of neural activity?

4. Under which operating conditions does event detection remain reliable as signal-to-noise ratio, detection threshold, spike amplitude, firing rate, and temporal firing statistics vary?

Chapter 3 addresses these questions through preliminary single-channel biological recordings and controlled synthetic-signal tests of the 15-channel MCU and 60-channel FPGA implementations. Evaluation includes analog performance, output data rate, detection precision and recall, temporal accuracy, and latency.

Chapter 3

Event-Driven Neural Acquisition Platform

3.1 Introduction and Motivation

As discussed in the previous chapter, conventional MEA acquisition systems generally preserve the complete extracellular waveform by continuously sampling every recording channel at frequencies sufficient to reconstruct neural spikes. This approach is essential when the experimental objective requires waveform morphology, waveform-based spike sorting, unit discrimination, artifact inspection, or retrospective reprocessing of the original voltage traces. However, many *in vitro* applications rely predominantly on the occurrence and spatial distribution of neural activity. Examples include firing-rate and burst analysis, network synchronization and propagation studies, long-term monitoring of network development, pharmacological and neurotoxicity screening based on activity patterns, and real-time closed-loop neuromodulation. In these cases, continuous full-waveform acquisition can be over-dimensioned with respect to the information ultimately used. The resulting data throughput, communication requirements, storage burden, and system cost all become increasingly restrictive as the number of recording channels and the duration of the experiment increase.

The platform presented in this chapter adopts a different acquisition paradigm. Extracellular spikes are detected directly in hardware and neural activity is represented as a stream of events encoding the corresponding channel and temporal information.

Within application domains in which the relevant biological information is the spatiotemporal organization of neural firing, this event-driven representation is intended as an alternative to general-purpose continuous-waveform MEA systems. The principal limitation is the loss of the original spike morphology, which prevents waveform-based sorting and makes the approach unsuitable when detailed signal reconstruction is required. In return, the event-driven paradigm provides ultra-low data throughput and storage, supports lower-cost implementations, and facilitates scaling toward larger channel counts. Hardware-level detection can also provide high temporal resolution and immediate access to multichannel activity maps, making the architecture particularly suitable for real-time raster visualization, network-level monitoring, and closed-loop applications in which detected events must rapidly trigger an external action.

The proposed system is a low-cost, modular, and ultra-low data-rate neural acquisition platform for 60-channel MEAs, developed through successive implementation stages. A single-channel AFE was initially designed as a demonstrator for the amplification and filtering of microvolt-level extracellular activity. Then, a multichannel motherboard was developed as the hardware core of the platform. The motherboard provides the mechanical and electrical interface to standard passive 60-channel MEAs and AFE modules, distributes power supply, and implements the event-driven acquisition core through threshold-comparator stages. The resulting outputs are processed by a control module. Two control implementations were developed: a 15-channel MCU-based version used to establish and validate the complete event-driven acquisition workflow, and a 60-channel FPGA-based version introduced to improve scalability, parallel recording and timing performance [6].

This chapter addresses the research questions introduced in Section 2.7.2 through the development and staged validation of the proposed platform. The experimental results characterize its performance and operating limits, providing a basis for discussing its suitability for applications centered on spatiotemporal neural activity.

This chapter first defines the design strategy and presents the common hardware architecture of the platform. The 15-channel MCU-based implementation is then described as the initial system-level proof of concept, followed by the functional architecture of the 60-channel FPGA-based implementation. The subsequent sections present the assessment of the analog front end performance, the controlled validation protocols used for the two backends, and the corresponding results.

3.2 Design Strategy and System Architecture

3.2.1 Event-Driven Neural Acquisition Requirements

The platform requirements were derived from the decision to preserve the spatiotemporal organization of neural events instead of the complete extracellular waveform. The analog front end therefore had to read out microvolt-level extracellular signals, reject electrode offset and slow baseline drift, attenuate out-of-band interference, and provide sufficient gain for reliable threshold comparison. At the output of the analog chain, spikes from each channel had to be detected and driven to the control unit in a simplified form. The digital backend had to detect and timestamp the event signal with sub-millisecond temporal precision, preserving the identity of the originating channel. Simultaneous events on different channels had to be captured without channel-dependent prioritization, serial polling, or software-service delays that could introduce event loss or timing skew as the channel count increased. The detected activity also had to be organized into a compact event stream for transmission to a host software. At host level, the software had to reconstruct the time–channel event representation, display the spatial distribution of neural activity in real time through a raster plot, and store the acquired event stream for subsequent analysis.

The principal requirements and the corresponding architectural choices are summarized in Table 3.1.

3.2.2 Event-Driven Acquisition Principle

As discussed in Section 2.2, spike amplitude and morphology depend on the electrode–cell coupling and on the acquisition chain. For network-level analyses, however, each detected spike can be treated as a discrete occurrence associated with a recording site and an occurrence time. Conventional MEA systems preserve the complete waveform by sampling every channel continuously, typically at 10 kHz to 50 kHz per channel [8, 21, 23]. By contrast, spontaneous per-electrode firing in dissociated neuronal cultures is commonly observed from sub-hertz values to a few tens of events per second, with substantially higher instantaneous activity during network bursts [15, 37]. Because an extracellular spike occupies only a few milliseconds, most

Table 3.1 Principal design requirements and corresponding architectural choices of the event-driven neural acquisition platform.

Requirement	**Target or design choice**	**Architectural implication**
MEA compatibility	Standard passive arrays with up to 60 recording channels	Spring-contact MEA interface and modular per-channel AFE connection
Extracellular-signal acquisition	Microvolt-level input and spike-band conditioning	High-input-impedance, low-noise amplification with offset rejection and band-pass filtering
Hardware event generation	Binary representation of threshold-crossing activity	Per-channel bipolar threshold detection with hysteresis
Channel variability	Independently calibrated detection thresholds	Per-channel calibration traces, noise estimation, and DAC-programmed comparator references
Event timing and concurrency	Sub-millisecond timing and capture of simultaneous channel activity	Hardware edge detection referenced to a common time base; parallel event monitoring in the scalable implementation
Event representation	Preservation of channel identity and event timing	Compact time–channel masks, timestamp association, buffering, and packet assembly
Host interaction	Real-time supervision, visualization, and storage	Bidirectional host communication, raster-plot reconstruction, and event-data logging
Scalability and accessibility	15-channel proof of concept followed by 60-channel operation	Replaceable MCU or FPGA control modules, commercially available components, and modular boards

continuously acquired samples in a spike-oriented recording describe inter-event baseline fluctuations and noise rather than neural events of interest.

The event-driven approach exploits this temporal sparsity by extracting spike occurrences before continuous multichannel digitization and transmission. The threshold-based detection principle was introduced in Section 2.7.1. In the proposed platform, the conditioned analog signal of each channel is compared with an independently programmed threshold. A threshold crossing causes the corresponding comparator output to change from the inactive logic state (0) to a logic-high pulse (1), thereby converting the analog spike into a digital event. The control unit detects the low-to-high transition, retains the identity of the originating electrode, and associates the event with the common system time reference. The resulting neural activity can be expressed conceptually as

$$\mathcal{E} = \{(i_k, t_k)\}_{k=1}^{N_{\mathrm{ev}}}, \tag{3.1}$$

Table 3.2 Main specifications of the event-driven neural acquisition platform. The MCU and FPGA entries describe successive control-unit implementations.

Specification	Value
MEA compatibility	Standard passive 60-channel MEAs
Extracellular spikes	10 µV to 500 µV amplitude; 0.1 kHz to 5 kHz bandwidth; 0 Hz to 300 Hz firing rate
Analog front end	One plug-in module per channel; nominal gain ≈ 500; passband ≈ 0.1 kHz to 5 kHz; input-referred noise < 10 µV RMS; output voltage range < 2.5 V
Overall analog gain	≈ 2500, comprising the AFE module gain (≈ 500) and the subsequent motherboard gain (5).
Hardware event detection	Per-channel Schmitt-trigger comparator with independently programmable threshold
Calibration resources	60 x 12-bit ADC (baseline acquisition); 60 x 8-bit DAC (threshold generation)
MCU-based backend	STM32H723ZG, Arm Cortex-M7 at 550 MHz; 15 channels; hardware input capture with 10 µs nominal resolution
MCU event representation	16-bit mask encoding 15-channel activity in 1 ms time slot
MCU communication	USB; approximately 2 kB s^{-1}; 2 packets per second comprising a 3-byte time header and 1000-byte array (16-bit activity mask in 500 ms);
FPGA-based backend	Lattice MachXO2-7000HE at 50 MHz; 60 event channels monitored in parallel
FPGA event representation	1 ms time slot; 96-bit word comprising a 32-bit timestamp, 60-bit channel-activity mask, and 4-bit parity field
FPGA communication	UART; approximately 12 kB s^{-1}
Host software	MATLAB graphical interface for calibration, acquisition control, real-time raster visualization, and event logging
Prototype cost	Less than USD 3000 for the complete 60-channel configuration. See section 3.7.5 for details.

where i_k identifies the recording channel and t_k is the assigned event time. The exact packet format differs between the MCU- and FPGA-based implementations, but both retain the same essential information: which electrode became active and when.

3.2.3 System Overview

The complete platform is organized into five functional blocks: the passive MEA interface, the plug-in analog front ends, the motherboard, the digital control unit, and the host computer. Figure 3.1a shows the common signal path and the two successive implementations of the control unit.

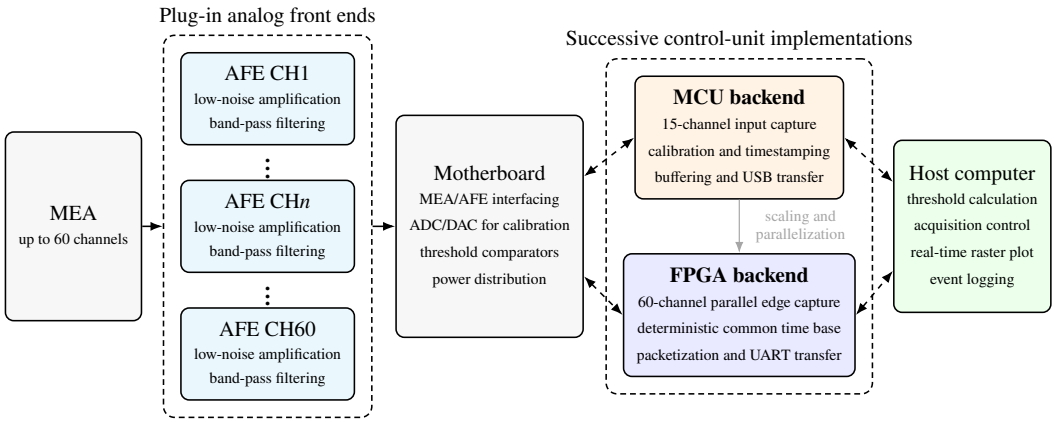
The platform interfaces with standard passive MEAs through spring-loaded contacts, which connect the array pads to the motherboard while allowing the biological preparation to be replaced without permanent wiring. Each electrode signal is routed to a vertical plug-in AFE module. The AFE provides low-noise amplification and band-pass conditioning before returning the signal to the motherboard.

The motherboard forms the electronic core of the system. It provides the MEA and AFE interfaces, regulated power distribution, and additional gain and level-shifting stage used to adapt each conditioned signal to the event-detection circuitry. Per-channel Schmitt-trigger comparators detect spikes and produce binary event lines. The comparator thresholds are individually programmable, allowing them to be adapted to the noise level of each channel. During calibration, on-board 12-bit analog-to-digital converters (ADC) acquire the conditioned analog signals, while 8-bit digital-to-analog converters (DAC) generate the corresponding comparator-reference voltages.

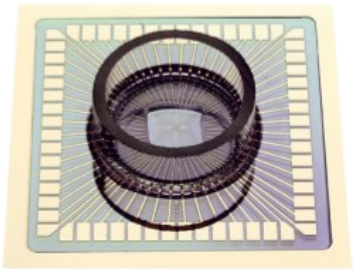
Two successive control modules were developed around this shared hardware. The first version used an STM32H723ZG microcontroller and operated on 15 event channels, responsible of threshold calibration, event detection, buffering and universal serial bus (USB) communication with the host computer [38]. This implementation established the first complete acquisition workflow and enabled rapid firmware-level development. The second version replaced the MCU with a Lattice MachXO2-7000HE FPGA and extended the backend to all 60 channels [39]. The FPGA monitors the comparator outputs, detects events using a deterministic common time base, supervises calibration, assembles timestamped activity masks, and streams the buffered data to the host through a universal asynchronous receiver-transmitter (UART) interface.

The host software completes the acquisition chain, triggering calibration and acquisition phases, decoding the event packets, reconstructing the multichannel activity map in real time, and storing the event stream for subsequent analysis. Depending on the specific control unit backend, two slightly different software versions were implemented.

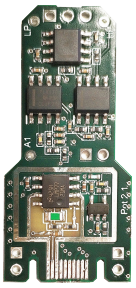
The principal specifications of the complete platform and of the two digital backends are summarized in Table 3.2.



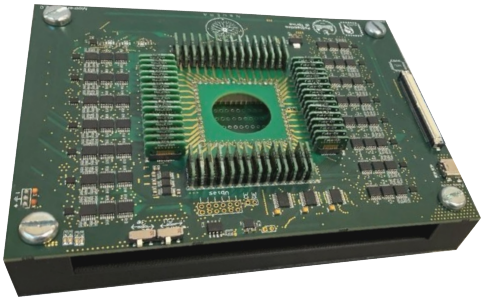
(a) High-level system architecture



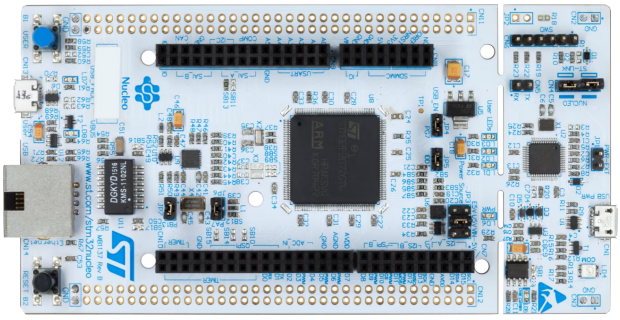
(b) Standard passive 60-channel MEA



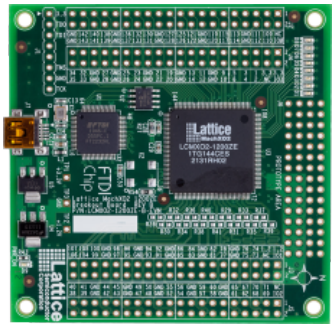
(c) AFE module



(d) Motherboard module



(e) STM32 Nucleo-144 board.



(f) Lattice MachXO2 Breakout Board.

Fig. 3.1 Overview of the event-driven neural acquisition platform. (a) High-level block diagram of the complete system, including the passive MEA, the modular analog front-end, the motherboard implementing signal interfacing and hardware spike detection, the digital control unit, and the host computer for visualization and storage. (b) Standard MCS 60-electrode MEA footprint [1]; (c) Photograph of the modular analog front-end boards. (d) Photograph of the motherboard hosting the MEA interface and event-detection circuitry. (e) STM32 NUCLEO-H723ZG development board [2]; (f) Lattice MachXO2 Breakout Board [3].

3.3 Shared Hardware Platform

3.3.1 Modular Organization

The acquisition platform comprised multiple printed circuit board (PCB) modules and was designed around the pad layout of standard passive 60-channel Microelectrode Arrays (MEAs) from Multi Channel Systems (MCS), as shown in Figure 3.1b, with 59 recording electrodes and one reference electrode [1].

The motherboard therefore integrates an array of spring-loaded contacts positioned to match the peripheral MEA pads. A dedicated 3D-printed holder aligns the array with the contact system and mechanically secures it to the motherboard through a screw-based fixture. This arrangement provides stable electrical contact while allowing the biological preparation to be removed and replaced without permanent bonding. A circular opening in the central region of the motherboard leaves the culture chamber accessible from above, thereby permitting optical inspection and the administration of pharmacological compounds during an experiment.

Each electrode signal is routed from the MEA interface to a dedicated plug-in AFE board. The 60 AFE modules are inserted vertically into edge connectors mounted on the motherboard. This organization separates the most noise-sensitive analog circuitry from the motherboard-level event-detection and control electronics, while allowing individual channels to be assembled, tested, replaced, or upgraded independently. The vertical arrangement also limits the planar footprint of the platform by exploiting the volume above the motherboard PCB.

The motherboard is connected to the digital control unit through a flexible flat cable (FFC), which carries the digital event outputs, calibration interfaces, and system-control signals. Because the two platform versions employ different control boards, two dedicated adapter PCBs were developed. Each adapter includes the same FFC interface used on the motherboard and provides the mechanical and electrical connections required by the corresponding control unit: the STM32 Nucleo-144 board in the MCU-based implementation and the Lattice MachXO2 Breakout Board in the FPGA-based implementation [2, 3]. This interface strategy allowed the digital backend to be replaced without redesigning the MEA interface, AFE modules, or motherboard.

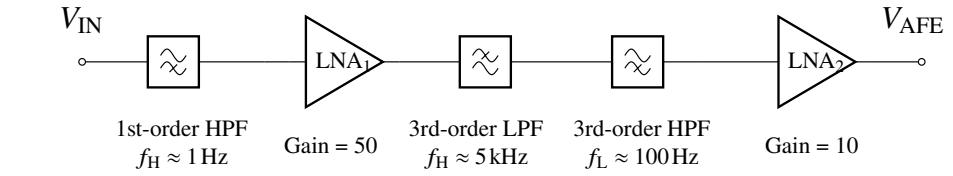
3.3.2 Low-Noise Analog Front End

The analog front end was initially developed as a single-channel potentiometric recorder and was subsequently adopted as the modular input stage of the multichannel platform. A photograph of the final boards is shown in Figure 3.1c. Its function is to convert the microvolt-level extracellular signal acquired from each MEA electrode into a sufficiently amplified and band-limited voltage for subsequent processing.

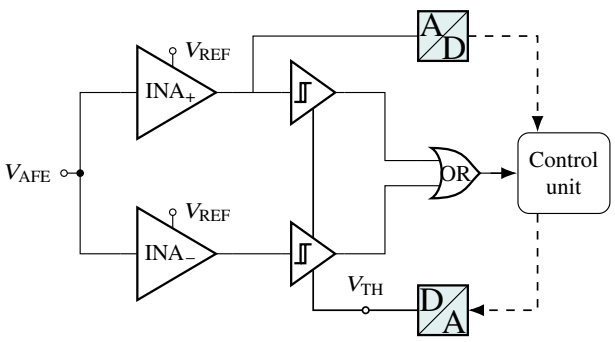
The design follows a cascaded amplification and filtering strategy rather than concentrating the entire gain in a single stage. At the input, a passive high-pass network attenuates the quasi-static electrode offset before the first active stage. The first amplifier is based on the LMP7721 precision operational amplifier, selected for its extremely low input-bias current, low input-voltage-noise density, and high input impedance [40]. Particular attention was given to the PCB layout around the input node. Short traces and a driven guard were used to limit leakage currents and parasitic coupling in this high-impedance region. Following the first low-noise amplification stage, a third-order Sallen–Key low-pass filter and a third-order Sallen–Key high-pass filter define the frequency range retained for extracellular spike acquisition. A second amplification stage then raises the signal amplitude before it is returned to the motherboard.

The nominal overall gain of the AFE is approximately 500, while its bandwidth extends from approximately 100 Hz to 5 kHz. The motherboard subsequently provides a further gain of 5, resulting in a nominal analog gain of approximately 2500 before threshold detection. This value was selected to acquire input-referred signals with amplitudes up to 500 µV without causing saturation along the analog chain. A higher AFE gain of approximately 2000 was instead adopted during the standalone AFE characterization described in Section 3.6.2. In those experiments, the motherboard gain stage was not included.

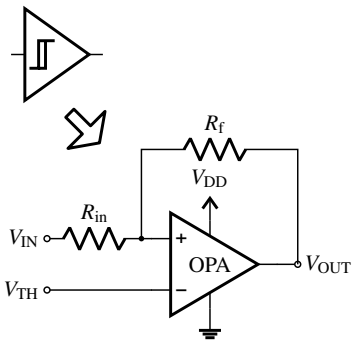
Following the LMP7721 specifications, the AFE was supplied by symmetrical ±2.5 V rails [40]. Because the electrode signal has microvolt-level amplitudes, it is particularly sensitive to noise, ripple, or drift of the supply rails. The supply for the analog chain was consequently generated using dedicated low-noise linear regulators located on the motherboard, as described in Section 3.3.3.



(a) Modular low-noise AFE signal path.



(b) Motherboard event-detection path.



(c) Non-inverting Schmitt-trigger.

Fig. 3.2 Single-channel diagram of the event-driven acquisition platform: (a) simplified AFE diagram, showing the multiple amplifying and filtering stages for signal readout; (b) motherboard event-detection diagram, consisting of two polarity paths for detecting both positive and negative spikes; and (c) non-inverting Schmitt trigger topology used for hardware threshold detection.

3.3.3 Motherboard

The motherboard completes the analog conditioning chain, provides the hardware resources required for channel-specific calibration, and converts each continuous AFE output into a single digital event signal.

After returning from the AFE, each conditioned signal is split into two parallel instrumentation-amplifier (INA) paths configured with opposite polarity. Both INAs provide a gain of 5 and reference their outputs to $V_{REF} = 1.65\,V$, generated by a dedicated voltage-reference circuit. This level shift centers the amplified signals within the unipolar 0 V to 3.3 V operating range of the calibration and threshold-detection circuitry.

The two-path architecture enables the detection of both positive- and negative-going extracellular spikes. The non-inverting branch preserves positive deflections, whereas the inverting branch converts negative deflections into positive excursions with respect to V_{REF}. Event detection is therefore independent of the polarity

introduced by the extracellular waveform, electrode coupling, or preceding analog stages. Each INA output is connected to an independent Schmitt-trigger stage. A threshold crossing in either branch produces a digital pulse, and the two trigger outputs are combined through a digital OR gate.

The hardware threshold detectors are implemented as non-inverting Schmitt triggers based on the TLV8544 operational amplifier [41] and supplied from a unipolar 3.3 V rail. When the conditioned signal exceeds the upper switching threshold, the corresponding output transitions to the logic-high state and remains high until the signal falls below the lower switching threshold. Positive feedback therefore introduces a finite hysteresis interval that suppresses repeated transitions when the signal or background noise remains close to the programmed threshold.

For the adopted configuration, the input-referred hysteresis width is

$$V_{\mathrm{H}} = \frac{R_{\mathrm{in}}}{R_{\mathrm{f}}} \left(V_{\mathrm{OH}} - V_{\mathrm{OL}}\right), \tag{3.2}$$

where R_{in} and R_{f} are the input and positive-feedback resistors, respectively. With $R_{\mathrm{in}} = 10\,\mathrm{k\Omega}$, $R_{\mathrm{f}} = 330\,\mathrm{k\Omega}$, and a nominal output swing from 0 V to 3.3 V, the resulting hysteresis is approximately 100 mV.

The TLV8544 was selected primarily for its very low quiescent current, which is advantageous because the threshold-detection circuit is replicated across all recording channels. However, its limited gain bandwidth, slew rate, and large-signal recovery characteristics constrain the temporal response of the event-generation stage [41]. This component therefore represents the principal hardware limitation affecting the end-to-end spike detection latency, as discussed in Section 3.8.

The motherboard also integrates onboard ADCs and DACs required for channel-specific threshold calibration. During the calibration phase, each conditioned signal, at the output of its non-inverting INA, is acquired by a dedicated 12-bit ADC. The samples are transferred via Serial Peripheral Interface (SPI) to the control unit, which estimates the channel noise level and calculates the corresponding detection threshold. The resulting value is then transmitted via SPI to the dedicated 8-bit DAC associated with that channel, whose output provides the reference voltage for both Schmitt-trigger branches. This process is repeated sequentially for each channel. The threshold-estimation procedures used by the MCU- and FPGA-based implementations are described in Sections 3.4.2 and 3.5.2, respectively.

Continuous waveform digitization is required only during calibration. During normal event-driven acquisition, the ADC path is inactive and the control unit receives only the digital event signals generated by the Schmitt-trigger and OR-gate network.

The analog circuitry is powered through an external bipolar-supply connector accepting nominal ±3.6 V to ±20 V. These rails can be provided either by a laboratory power supply or by batteries. The positive rail is regulated to 2.5 V by an LT3042 ultralow-noise positive low-dropout regulator, whereas the negative rail is regulated to −2.5 V by an LT3093 ultralow-noise negative regulator [42, 43]. The resulting low-noise bipolar supplies power the AFE modules and the precision analog circuitry on the motherboard. The digital control unit is instead powered through the USB connection to the host computer, separating the sensitive analog supply domain from the MCU- or FPGA-based digital domain.

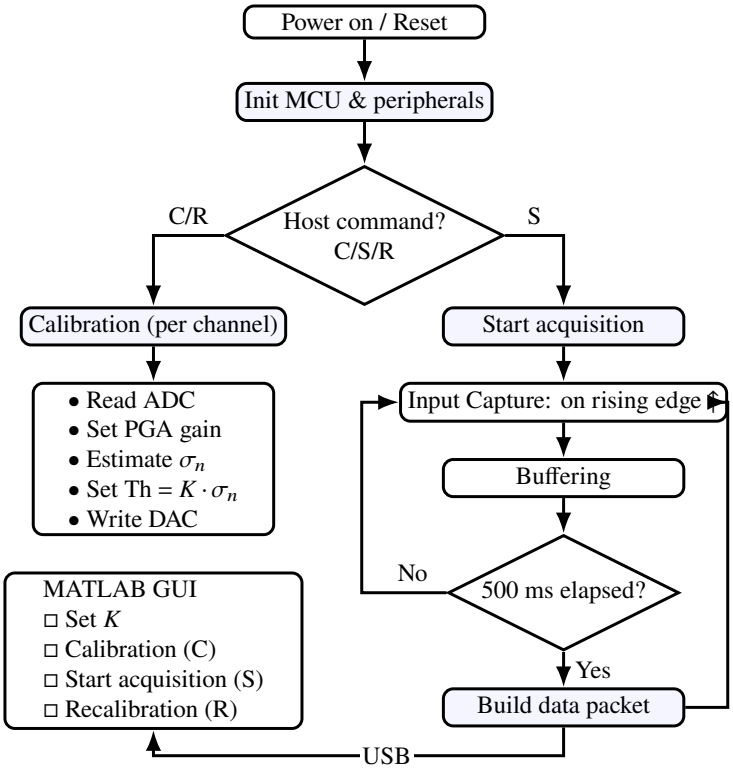
3.4 MCU-Based 15-Channel Implementation

The first complete event-driven implementation was built around an STM32H7 microcontroller and enabled simultaneous monitoring of 15 active channels [6, 38]. This version was conceived as a full proof of concept of the event-driven acquisition chain, combining channel calibration, hardware-timed event capture, packet generation, USB communication, and host-side visualization in a single platform. The selected channel count was mainly driven by the available timer input-capture resources of the microcontroller and by the need to keep the first implementation relatively simple and experimentally accessible.

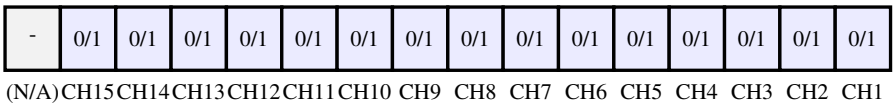
3.4.1 General Operating Sequence

The operating sequence was organized into two main phases. In the first phase, the system performs channel-by-channel calibration by sampling the analog signals sequentially, estimating the background noise, and programming the corresponding detection thresholds. In the second, the MCU acquires only digital event outputs generated by the motherboard.

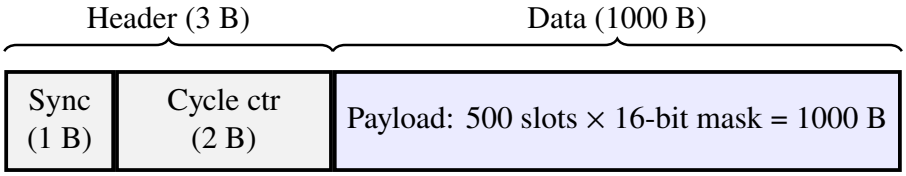
After power-up, the firmware initializes the timer peripherals, ADC and DAC interfaces, USB communication, and the event buffers. The host computer then sends



(a) Firmware operating sequence and host interaction.



(b) Encoding of one 1 ms activity word.



(c) Packet organization used for USB transmission.

Fig. 3.3 Operation and data formatting of the MCU-based 15-channel implementation. (a) Simplified workflow of initialization, calibration, acquisition, and host interaction. (b) Binary activity word representing one 1 ms time bin. (c) Packet structure used to transmit 500 ms blocks of event activity to the host computer.

the command that determines the operating mode: calibration, start of acquisition, or recalibration. The firmware workflow is also summarized in Figure 3.3a.

3.4.2 Calibration Phase

The firmware performs calibration sequentially, one channel at a time. For each channel, the conditioned analog signal is sampled for 1 s at $10\,\mathrm{kSs}^{-1}$, yielding $N = 10000$ samples from which the background-noise level is estimated.

First, the calibration waveform is centered by subtracting its sample mean:

$$x_{c,i}[n] = x_i[n] - \bar{x}_i, \qquad \bar{x}_i = \frac{1}{N}\sum_{n=1}^{N} x_i[n], \tag{3.3}$$

where $x_i[n]$ denotes the acquired calibration samples of channel i and $x_{c,i}[n]$ denotes the corresponding zero-mean centered samples.

The noise level is then estimated using the median-based estimator proposed by Quiroga et al. for extracellular spike detection [30]:

$$\hat{\sigma}_{n,i} = \frac{\mathrm{median}_n\left(|x_{c,i}[n]|\right)}{0.6745}, \tag{3.4}$$

$$T_i = K\hat{\sigma}_{n,i}, \tag{3.5}$$

where $\hat{\sigma}_{n,i}$ is the estimated noise standard deviation, T_i is the corresponding threshold amplitude, and K is a configurable scaling coefficient. For zero-mean Gaussian noise, the factor 0.6745 converts the median absolute amplitude into an estimate of the standard deviation.

The median-based amplitude estimator reduces sensitivity to sparse spikes and artifacts. Nevertheless, these transients can affect the sample-mean baseline estimate, and high firing activity during calibration can increase the estimated noise level and the resulting detection threshold.

The MCU-based implementation uses $K = 5$ by default. This conservative setting reduces false detections caused by background fluctuations, but can also decrease sensitivity to spikes whose amplitudes approach the noise floor. Lower values increase sensitivity at the expense of a higher false-positive rate. For this implementation, an analysis of the effect of K on the performance was not performed.

Once T_i has been calculated, it is mapped to the corresponding 8-bit DAC code and programmed on the motherboard through SPI. The DAC output defines the reference voltage applied to both Schmitt-trigger branches of the selected channel.

The procedure is repeated for all 15 active channels before event detection begins. Independent calibration partially compensates for channel-to-channel differences in electrode impedance, cell–electrode coupling, analog offset, and noise amplitude.

3.4.3 Event-Driven Acquisition Phase

Each motherboard event line is connected to a timer input-capture channel. During event-driven acquisition, a rising edge causes the corresponding timer value to be latched in hardware before execution of the interrupt service routine. This mechanism prevents variable interrupt latency from directly affecting the captured timestamp and provides a nominal internal temporal resolution of 10 µs. All capture timers share a common timing reference, ensuring consistent timestamp alignment across channels.

A refractory interval of 2 ms is enforced independently on each channel to suppress repeated detections caused by the biphasic structure of an extracellular spike or by residual switching around the threshold. The refractory interval imposes a theoretical maximum accepted rate of 500 Hz per channel. This ceiling was considered adequate for the proof-of-concept implementation and for the network-level activity targeted by the platform. It should nevertheless be interpreted as an architectural operating compromise rather than as a universal upper bound on neural firing activity.

For transmission, detected events are represented in 1 ms time bins. Each bin is encoded as a 16-bit activity word in which 15 bits correspond to the active channels and one bit remains unused, as shown in Figure 3.3b. A bit is asserted whenever an event occurs on the corresponding channel during that time window.

A block of 500 consecutive activity words represents 500 ms of acquisition and forms the payload of one transmission packet, as shown in Figure 3.3c. Each packet contains a 3-byte header, comprising one synchronization byte and one 16-bit cycle counter, followed by 500 activity words. Since two packets are transmitted per second, the nominal data rate is fixed to approximately $2\,\mathrm{kB\,s^{-1}}$ and does not increase with the number of detected events.

3.4.4 Host Interface

A host application manages command transmission, calibration control, packet decoding, real-time visualization, and data storage. Through the graphical interface,

the operator can configure the threshold multiplier, start calibration or acquisition, stop the recording, and request recalibration when required.

The received USB packets are decoded directly into a time–channel activity matrix. The resulting activity is displayed as a real-time raster plot and stored for offline analysis or comparison with controlled synthetic stimuli.

3.4.5 Architectural Limitations

The MCU-based implementation demonstrated that automatic threshold calibration, hardware-timed event capture, low-data-rate packetization, USB communication, and real-time raster visualization could be integrated into a single proof-of-concept platform. Its principal limitations were architectural rather than conceptual.

The number of simultaneously monitored channels was first constrained by the available timer input-capture peripherals and by their mapping to the MCU pins. Although each timer captures the event timestamp in hardware, subsequent operations remain firmware-driven: the interrupt service routine identifies the channel, enforces the refractory interval, and updates the activity mask. These operations are executed sequentially. Increasing the number of channels and the occurrence of simultaneous or closely spaced events would therefore increase the worst-case interrupt load and make bounded processing latency more difficult to guarantee.

The transmitted representation also preserves a coarser temporal resolution than the timer hardware. Events are captured internally with 10 µs resolution but reported through 1 ms activity bins, while the 2 ms refractory interval introduces a per-channel dead time and limits the accepted event rate to 500 Hz.

The MCU version was therefore an effective demonstrator of the complete event-driven workflow, but it was not considered an appropriate architecture for extending the platform to all 60 channels. These constraints motivated the transition to an FPGA-based backend, in which all event inputs could be monitored concurrently using dedicated parallel logic and a common, deterministic time base.

3.5 FPGA-Based 60-Channel Implementation

The second control-unit implementation replaced the MCU with a Lattice MachXO2-7000HE FPGA operating at 50 MHz [39]. The shared AFE and motherboard hardware remained unchanged; the purpose of the new backend was to extend event acquisition to all 60 MEA channels and remove the sequential-processing constraints of the MCU implementation.

The FPGA coordinates motherboard calibration, monitors the 60 digital event lines concurrently, associates the detected activity with a common time reference, and transfers fixed-format event records to the host computer. The architecture is described here at the functional level.

3.5.1 General Operating Sequence

As in the MCU-based implementation, calibration and event acquisition are implemented as distinct operating phases. After initialization, the FPGA remains in an idle state while waiting for a host command. The host can request calibration, start event-driven acquisition, stop the current recording, or repeat calibration when the recording conditions change. During calibration, the FPGA coordinates the sequential acquisition of the analog background signal from each channel and the subsequent programming of the channel thresholds. During event-driven acquisition, the ADC path is inactive and the FPGA continuously monitors the 60 binary event lines generated by the motherboard. The two phases are mutually exclusive. The complete operating sequence is summarized in Figure 3.4.

3.5.2 Calibration Phase

In the FPGA-based implementation, calibration is divided between deterministic hardware coordination and host-side statistical processing. The FPGA does not estimate the noise level or calculate the detection thresholds. Its role is to control the motherboard ADCs and DACs, and transfer the calibration data to the host software.

When the host sends the calibration command, the FPGA enters calibration mode and sequentially interrogates the 60 motherboard ADC channels. For each channel, it acquires 1000 samples at $10\,\mathrm{kS\,s^{-1}}$, corresponding to a 100 ms observation interval.

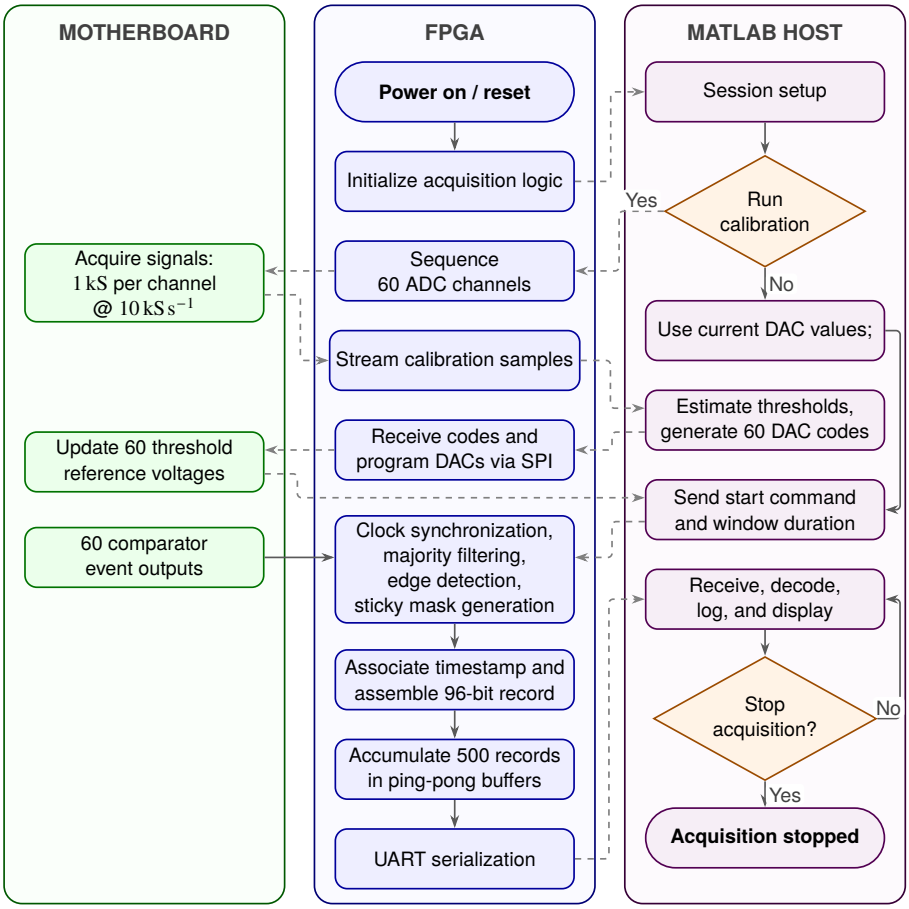
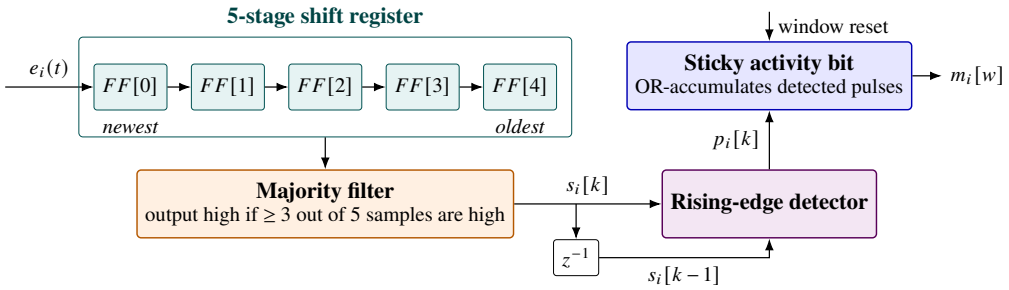


Fig. 3.4 Operating sequence of the FPGA-based implementation, showing the distribution of initialization, optional calibration, parallel event acquisition, buffering, communication, and host-side processing across the motherboard, FPGA, and MATLAB host.

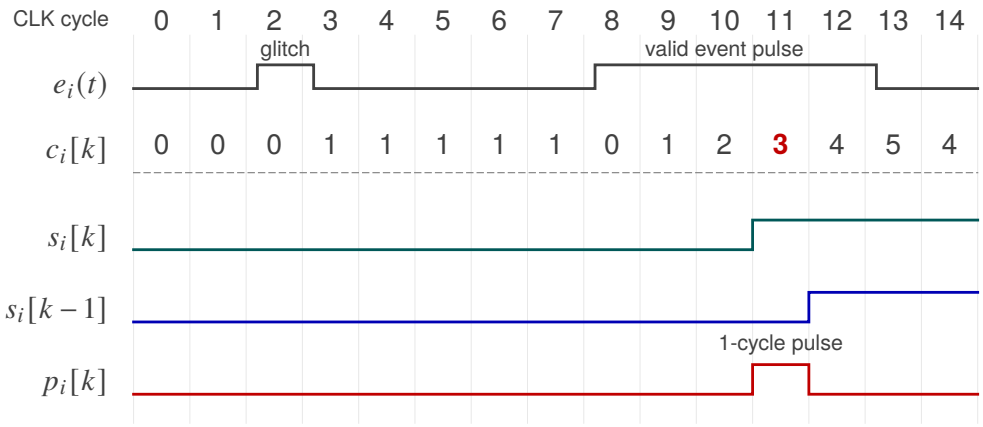
The FPGA transfers the resulting 12-bit samples to the host without performing numerical processing. Each sample is transmitted as two bytes, resulting in a total transfer of 120 000 bytes. Sequential sampling of all channels requires around 6 s.

After receiving the calibration dataset, the host reconstructs the channel waveforms and calculates one 8-bit threshold code per channel using the procedure described in Section 3.5.4. The resulting 60-byte table is returned to the FPGA in channel order.

The FPGA then programs the motherboard DACs sequentially through SPI, thus applying the threshold voltage to the Schmitt-trigger reference associated with a specific channel. Once all 60 outputs have been programmed, the FPGA reports completion and waits for the acquisition-start command. The programmed references remain active until they are overwritten or the motherboard is powered down.



(a) Per-channel event-conditioning path.



(b) Majority-filter and rising-edge timing.

Fig. 3.5 Parallel event conditioning in the FPGA backend. (a) Each asynchronous motherboard event line $e_i(t)$ is sampled through a five-stage register. A three-out-of-five majority rule produces the filtered level $s_i[k]$, while a one-cycle delay provides $s_i[k-1]$. Their combination, $p_i[k] = s_i[k] \wedge \neg s_i[k-1]$, generates one pulse on each accepted rising transition. This pulse is OR-accumulated into the sticky occupancy bit $m_i[w]$, which is cleared after the current acquisition window has been captured. The datapath is replicated independently for all 60 channels. (b) An isolated transient produces only one high sample and is rejected, whereas a sustained event raises the five-sample population count $c_i[k]$ to three, asserts $s_i[k]$, and generates a single-cycle rising-edge pulse.

This partition assigns timing-critical peripheral control to the FPGA while retaining the threshold estimator in software. The estimator and its parameters can consequently be modified without changing the FPGA logic.

3.5.3 Parallel Event-Driven Acquisition

During acquisition, all 60 motherboard event outputs are monitored concurrently. Since the motherboard event outputs are asynchronous with respect to the 50 MHz

FPGA clock, each channel is processed by the identical conditioning path shown in Figure 3.5. The path transfers the event level into the FPGA clock domain, rejects isolated short transitions, and produces one clock-cycle pulse for each accepted rising edge. All 60 instances operate concurrently, without channel prioritization or interrupt-service delays.

The filtered level can be asserted after three out of five sampled-high states. At 50 MHz, this corresponds to 60 ns to 100 ns from the first captured high sample; because the external transition is asynchronous, its phase relative to the FPGA clock adds an uncertainty of up to one clock period (20 ns).

Each accepted rising edge sets the corresponding bit of a 60-bit sticky activity register. Once asserted, a bit remains high until the end of the current acquisition window. The acquisition-window duration T_{w} is configurable and is set to 1 ms by default, similarly to the MCU implementation. At each window boundary, the complete 60-bit activity mask is captured and the sticky register is cleared for the next interval. Events occurring on different channels within the same window are retained as separate asserted bits. Multiple events from the same channel within one window are instead represented by a single bit; the transmitted quantity is therefore channel occupancy rather than an exact within-window event count.

Unlike the MCU implementation, the principal FPGA acquisition path does not enforce a per-channel refractory interval in hardware. A refractory rule can instead be applied consistently during host-side or offline processing.

At the end of each acquisition window, the 60-bit activity mask is associated with the common FPGA time base and assembled into a fixed-length 96-bit record. Each record contains a 32-bit timestamp, a 60-bit activity mask, and four parity bits. The implemented host software did not validate these bits; therefore, the prototype did not provide end-to-end packet-error detection.

Each record therefore occupies 12 bytes. With $T_{\mathrm{w}} = 1\,\mathrm{ms}$, the FPGA generates one record per millisecond and accumulates 500 consecutive records in one memory bank. Each bank therefore contains 6000 bytes and represents 500 ms of activity. Two banks operate in ping-pong mode: while one is filled, the other is transmitted through the $2\,\mathrm{Mbit\,s^{-1}}$ UART interface and the FTDI USB bridge. The resulting average payload rate is approximately $12\,\mathrm{kB\,s^{-1}}$.

The timestamp counter has a nominal resolution of 1 µs and wraps after approximately 71.6 min; rollover is handled during host-side reconstruction. Nevertheless, all events accumulated within the same window share one timestamp. The effective temporal resolution of the transmitted activity is therefore determined primarily by the window duration T_w, rather than by the finer counter resolution.

3.5.4 Host Software and Threshold Estimation

A MATLAB host application supervises configuration, threshold calibration, acquisition control, packet decoding, visualization, and data storage. Before a session, the operator sets acquisition-window duration T_w, threshold multiplier K, optional post-processing refractory interval, and the subset of channels to be displayed. The default acquisition window is 1 ms, while the default refractory interval is 2 ms and can be disabled when required by the experimental protocol.

When calibration is selected, the host sends the corresponding command to the FPGA, initiating the hardware sequence described in Section 3.5.2. The host therefore receives 1000 samples per channel, corresponding to a time calibration window of 100 ms per channel.

To obtain a suitable threshold level, a two-stage robust estimator is applied independently to every channel. The procedure is designed to prevent spontaneous spikes occurring during the calibration window from being interpreted as background noise. Quiroga median-based estimator alone demonstrated robustness to spike occurrence in case of low neural activity. However, if a higher-frequency burst occurs during the 100 ms calibration window, the median-based estimator could increase the estimated noise level and produce an unnecessarily high detection threshold.

First, the channel baseline is removed from the signals as

$$x_{c,i}[n] = x_i[n] - \text{median}_n(x_i[n]), \tag{3.6}$$

where $x_i[n]$ is the calibration waveform for channel i. A preliminary noise estimation is then obtained from the 25th percentile of the absolute centered samples:

$$\widehat{\sigma}_{P25,i} = \frac{P_{25}\left(\left|x_{c,i}[n]\right|\right)}{0.3186}. \tag{3.7}$$

The use of this lower percentile makes the preliminary estimate less sensitive to sparse, high-amplitude transients than a median-based estimator. Samples satisfying $\left|x_{c,i}[n]\right| > 4\widehat{\sigma}_{P25,i}$ are identified as potentially contaminated by spikes. These samples are excluded from the signal for the final estimation, using an exclusion mask. The exclusion mask is extended by 1.5 ms before and after the flagged index, removing the neighboring samples belonging to the same transient rather than excluding only its peak. The samples not covered by the mask form the clean set $x_{\text{clean},i}$. If at least 200 samples, corresponding to 20% of the original calibration window, remain available, the final noise estimate is calculated using the robust median-absolute-deviation (MAD) estimator

$$\widehat{\sigma}_{\text{clean},i} = \frac{\text{median}\left(\left|x_{\text{clean},i} - \text{median}\left(x_{\text{clean},i}\right)\right|\right)}{0.6745}. \tag{3.8}$$

When fewer than 200 clean samples remain, the percentile estimate is retained as a fallback. The noise estimate used for threshold generation is therefore

$$\widehat{\sigma}_i = \begin{cases} \widehat{\sigma}_{\text{clean},i}, & N_{\text{clean},i} \geq 200, \\ \widehat{\sigma}_{P25,i}, & N_{\text{clean},i} < 200. \end{cases} \tag{3.9}$$

The absolute threshold-reference voltage is obtained by adding a noise-dependent offset to the measured channel baseline:

$$V_{\text{TH},i} = \text{median}_n\left(x_i[n]\right) + \max\left(K\widehat{\sigma}_i, 10V_{\text{LSB}}\right), \tag{3.10}$$

where V_{LSB} is the voltage increment of the 8-bit threshold DAC. The lower bound of ten DAC increments prevents exceptionally quiet channels from receiving an unrealistically small threshold offset that could be lost through quantization or produce excessive event activity caused by small baseline oscillations. The threshold voltages for all the channels are finally quantized into an unsigned 8-bit DAC code and transmitted to the FPGA for DAC programming, as described in Section 3.5.2.

After threshold programming, the host sends the acquisition-start command, including the counter value corresponding to the selected accumulation window. During acquisition, incoming serial data are processed asynchronously in blocks of 600 bytes, while a separate 50 ms timer updates the graphical display. The software decodes the fixed-length records, reconstructs the 60-bit activity masks, checks

timestamp progression, handles counter rollover, and maps each bit to its physical MEA channel. The reconstructed activity is displayed as a real-time raster plot and stored in a fixed-format event log for subsequent analysis. Channel selection affects visualization and analysis but does not alter the parallel acquisition of all 60 hardware inputs. Optional operations, including a per-channel refractory filter, are applied after packet decoding and therefore do not modify the events captured by the FPGA.

3.5.5 Architectural Characteristics and Limitations

The FPGA backend resolves the principal scaling limitations of the MCU implementation. All 60 event inputs are observed concurrently, the event timestamping is referenced to a common deterministic time base, and the packet size and generation schedule remain independent of the number of simultaneously active electrodes. The architecture is therefore particularly suited to network bursts, in which several spatially distributed channels may become active within a short interval.

As described above, the configurable acquisition window determines the temporal granularity of the transmitted activity. With the default $T_{\mathrm{w}} = 1\,\mathrm{ms}$, multiple events occurring on the same channel within one window are represented by a single occupancy bit, despite the finer microsecond resolution of the internal timestamp counter. Reducing T_{w} would decrease the probability of within-channel event merging and improve temporal resolution, but would proportionally increase the packet-generation rate and output throughput.

The fixed-rate format also generates records for inactive windows. The throughput is therefore bounded and independent of neural activity, but it is not strictly proportional to the number of detected events. An address-event representation could reduce the data rate during sparse activity, although it would require variable-length packet handling.

3.6 Validation Methodology

3.6.1 Validation Strategy and Interpretation Boundaries

The platform was validated through a staged strategy reflecting its successive development phases. The single-channel AFE was first assessed using both microvolt-scale artificial signals and extracellular recordings from *in vitro* neuronal cultures to evaluate the electrode-interface and analog-conditioning stages without including hardware event detection. The MCU-based platform was subsequently evaluated as a 15-channel proof of concept of the complete event-driven workflow. Finally, the 60-channel FPGA-based backend underwent the most comprehensive controlled campaign. The MCU and FPGA campaigns bypassed the microvolt-level AFE and injected synthetic post-AFE signals directly into the motherboard event-detection path. In these backend experiments, the same waveform was distributed simultaneously to all tested channels. This configuration provided identical ground truth across the array and deliberately stressed the ability of the digital backend to capture concurrent events.

A complete end-to-end validation using neuronal cultures was not adopted because biological preparations are neither deterministic nor readily available for repeated engineering iterations and extensive parameter sweeps. Culture preparation and maturation require substantial time and resources. Moreover, living neurons do not provide the controlled event times and independently adjustable parameters required for quantitative detection metrics.

An alternative would have been to inject synthetic neural-scale signals directly at the AFE input. A precision source-measure unit was used for the preliminary single-channel AFE characterization, but was not available during the subsequent backend campaigns. More generally, generating microvolt-level extracellular-like waveforms with sufficiently low output noise, controlled amplitude, and representative source impedance requires specialized and costly instrumentation. The event-driven backends were therefore validated using millivolt-level signals representative of the AFE output, while the analog interface was characterized separately.

The resulting methodology defines clear interpretation boundaries. The AFE experiments demonstrate the feasibility of acquiring microvolt-level extracellular activity; the MCU campaign establishes the complete event-driven workflow as a

proof of concept; and the FPGA campaign characterizes the operating envelope of the scalable 60-channel backend under controlled conditions. Taken together, these results support the proposed architecture, but they do not constitute an end-to-end biological validation of the complete 60-channel system.

3.6.2 Preliminary Single-Channel AFE Demonstration

The AFE was assessed separately from the multichannel event-driven platform through two preliminary experiments. The first used controlled microvolt-scale artificial signals to verify the response of the analog chain. The second examined whether extracellular activity from a neuronal culture could be observed through the custom AFE when connected to a standard passive MEA. These experiments were intended as an initial demonstration of feasibility rather than as a complete validation of the analog front end.

Microvolt-Scale Electrical Test

Artificial pulse-like signals were generated using an Agilent B2912A Precision Source/Measure Unit [44]. The source was connected directly to the AFE input and configured to produce 1 ms pulses with amplitudes representative of extracellular recordings. A $50\,\mu V$ pulse was adopted as a representative test condition.

The SMU was operated in voltage-source mode and connected directly to the AFE input. The effective source impedance and the complete AFE input impedance were not independently characterized under these test conditions. This experiment therefore assessed the response to directly applied microvolt-scale voltage pulses. The effects of a representative electrode–electrolyte source impedance require separate characterization.

The source, AFE, and acquisition wiring were placed inside a Faraday enclosure to reduce environmental interference. Pulse-train measurements were used to inspect amplification, waveform conditioning, baseline recovery, and the visibility of the applied signal at the AFE output. Zero-input recordings were acquired with the AFE input shorted to ground to estimate the background noise of the electronic chain. The conditioned output was sampled at $10\,kS\,s^{-1}$ and transferred to a host computer for offline analysis.

Neuronal-Culture Demonstration

The biological demonstration was conducted in collaboration with the Department of Drug Science and Technology of the University of Turin. Primary midbrain dopaminergic cultures were prepared by the collaborating laboratory according to the protocol described in [45]. Cells were plated on a standard passive 60-channel MCS MEA and maintained until 14 days *in vitro*.

A custom adapter provided direct electrical access to selected MEA electrodes. The single-channel AFE was connected between one recording electrode and the array reference, and its conditioned output was digitized at $10\,\mathrm{kS\,s^{-1}}$. The connection was subsequently moved to a small number of additional electrodes, providing recordings from 4 different sites of the same preparation. Average acquisition duration per channel was approximately 80 s.

The recordings were analysed offline without additional digital filtering, so that the reported traces and noise values reflected the actual signal present at the AFE output. The noise floor was estimated as the root mean square (RMS) value of the recording after excluding time intervals containing burst activity, which would otherwise produce a substantial overestimation of the baseline noise. Sparse, low-rate spike-like events were not systematically removed and may therefore have introduced a limited positive bias in the RMS estimate. Candidate spike-like events were identified as local maxima exceeding a threshold equal to three times the estimated noise RMS, $T = 3\sigma_{\mathrm{noise}}$. All detected events were subsequently inspected visually to reject evident artifacts and confirm the presence of spike-like waveforms. Event amplitudes were measured relative to the local baseline, and the activity rate was calculated as the number of accepted events divided by the recording duration. Results are reported in Section 3.7.1.

The available dataset comprised only a limited number of recordings and electrodes and was not collected through a predefined repeated-measures protocol. Moreover, no simultaneous reference acquisition or recoverable commercial-system dataset was available for quantitative comparison. The experiment therefore does not establish channel yield, measurement accuracy, repeatability across cultures, long-term stability, or equivalence with established MEA instrumentation. A complete AFE validation would require repeated testing across multiple preparations

and electrodes, controlled electrical characterization of the electrode interface, and comparison with a calibrated reference acquisition system.

The culture preparation referenced above was conducted in accordance with European Directive 2010/63/EU and under the authorization reported in [45]. The engineering activity described here was limited to electronic recording of cultures prepared by the collaborating laboratory.

3.6.3 Common Synthetic-Signal Backend Test Bench

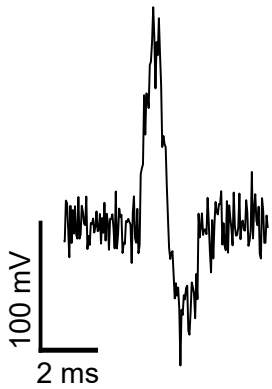
The MCU- and FPGA-based backends were assessed using the same general synthetic-signal test bench. The purpose was to characterize the motherboard-level event-detection path, the digital control unit, communication, and host-side event reconstruction under repeatable conditions with known event times.

The waveforms were generated in MATLAB and converted into an analog signal using a National Instruments USB-6259 data-acquisition device (DAQ) [46]. The DAQ analog output was updated at $25\,\mathrm{kS\,s^{-1}}$. A dedicated distribution PCB reproduced the peripheral contact layout of a standard 60-channel MCS MEA, allowing direct coupling to the spring-loaded contacts of the motherboard. The single DAQ output was distributed in parallel to the selected contact positions.

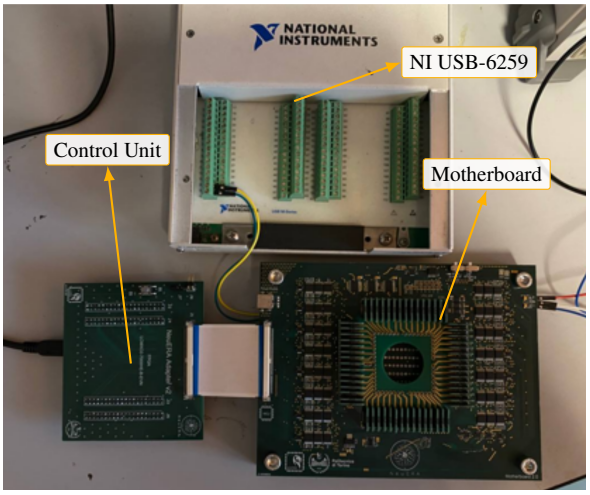
During these tests, the plug-in AFEs were replaced by passive bridge modules that routed the generated waveform directly to the motherboard INA inputs. The applied voltage was therefore selected to represent the post-AFE signal scale rather than the native microvolt-level extracellular signal.

From the injection point onward, the waveform followed the standard path shown in Figure 3.2b: amplification and level shifting, dual-polarity threshold detection, digital event generation, MCU- or FPGA-based processing, data transfer, and host-side reconstruction. The electrode interface and low-noise AFE were excluded from the tested path.

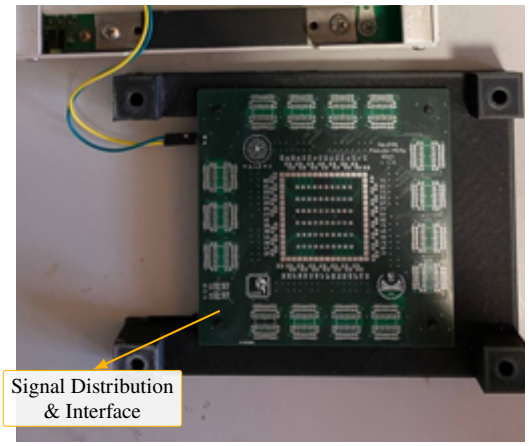
The same waveform was delivered simultaneously to all active channels. Each synthetic spike was therefore expected to produce concurrent threshold crossings across the tested population. This configuration provided identical ground truth for every channel and deliberately stressed the ability of the backend to process



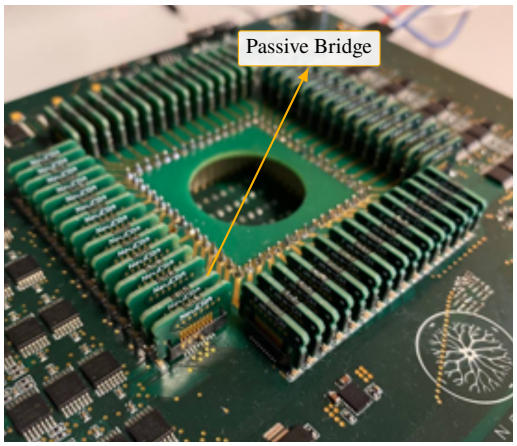
(a) Single synthetic spike waveform generated in MATLAB.



(b) Complete FPGA-based experimental setup.



(c) Signal-distribution PCB.



(d) Passive bridge modules for AFE bypassing.

Fig. 3.6 Synthetic-signal backend validation setup. (a) Single spike waveform generated in MATLAB and reproduced by the NI USB-6259. (b) Complete FPGA-based experimental setup, including the DAQ, FPGA control unit, and motherboard. (c) Signal-distribution PCB used to replicate the DAQ output across the 60-channel motherboard interface with the standard MEA footprint. (d) Passive bridge modules mounted in place of the analog front ends, allowing the synthetic post-AFE signal to be injected directly into the motherboard detection path.

simultaneous events. It did not reproduce the independent waveforms, channel-dependent noise, and partial synchrony normally encountered in a neuronal network.

Synthetic Signal Model

The dimensionless synthetic spike template was represented as the weighted difference of two time-shifted Gaussian components. Two-component Gaussian models have previously been used to parameterize recorded extracellular spike waveforms [47]:

$$s(t) = \exp\left[-\frac{(t-T/4)^2}{2\sigma_1^2}\right] - 0.5\exp\left[-\frac{(t-3T/4)^2}{2\sigma_2^2}\right], \quad t \in [0,T], \tag{3.11}$$

with

$$T = 2\,\text{ms}, \qquad \sigma_1 = 0.3\,\text{ms}, \qquad \sigma_2 = 0.4\,\text{ms}. \tag{3.12}$$

These parameter values define a fixed, asymmetric biphasic waveform on the millisecond timescale of extracellular spikes [8]. The Gaussian components are centred at $T/4 = 0.5\,\text{ms}$ and $3T/4 = 1.5\,\text{ms}$, respectively. Their widths and relative weight produce a narrower positive lobe followed by a broader, lower-amplitude negative lobe. Keeping the template fixed enables comparisons across test conditions without introducing waveform variability. A representative synthetic spike is shown in Figure 3.6a.

The complete stimulus was generated as

$$x(t) = A_{\text{spike}} \sum_k s(t - t_k) + \sigma_{\text{noise}} w(t), \tag{3.13}$$

where t_k is the onset time of the k-th spike, A_{spike} is its amplitude scaling parameter, σ_{noise} is the noise standard deviation, and $w(t) \sim \mathcal{N}(0,1)$ is a zero-mean, unit-variance white Gaussian process. The SNR was defined as the linear amplitude ratio

$$\text{SNR}_{\text{lin}} = \frac{A_{\text{spike}}}{\sigma_{\text{noise}}}. \tag{3.14}$$

In [6], the SNR values were incorrectly reported in decibels. In this chapter they are corrected and reported as the dimensionless linear amplitude ratio actually used to generate the stimuli.

White Gaussian noise was selected as a controlled benchmark for characterizing detection performance at defined SNR values. The synthetic signal model was

intentionally simplified: it did not reproduce electrode-dependent spike morphology, coloured or non-stationary biological noise, overlapping units, baseline drift, or channel-to-channel variability. The generated waveforms nevertheless retained a biphasic morphology, millisecond-scale duration, and a representative post-AFE amplitude scale, while allowing spike amplitude, noise, SNR, firing rate, and event timing to be controlled independently. The resulting measurements characterize detection performance for this specified morphology and noise model; generalization to more realistic recording conditions requires additional validation.

Common Detection Metrics

The known input times were compared with the reconstructed event stream. Because the two backends used different timing and alignment procedures, the specific matching rules are described in Sections 3.6.4 and 3.6.6. Once the detected and expected events had been classified, the common detection metrics were

$$\begin{gathered}\text{Precision} = \frac{TP}{TP+FP}, \qquad \text{Recall} = \frac{TP}{TP+FN}, \\ F_1 = 2\frac{\text{Precision}\,\text{Recall}}{\text{Precision}+\text{Recall}}.\end{gathered} \tag{3.15}$$

where TP denotes true positives, FP denotes false positives, and FN denotes false negatives.

The MCU campaign used regular spike trains and a simplified interval-based comparison to establish proof-of-concept operation. The FPGA campaign introduced explicit synchronization, clock-skew correction, one-to-one event matching, repeated sessions, and dedicated studies of threshold selection, amplitude variation, irregular timing, and propagation latency.

3.6.4 MCU Proof-of-Concept Protocol

The first backend campaign evaluated the 15-channel MCU implementation using the common test bench described in Section 3.6.3. Its objective was to demonstrate the complete event-driven workflow and identify broad performance trends as a function of signal quality and event rate.

The experimental design combined ten linear SNR values with five firing rates:

$$\mathrm{SNR}_{\mathrm{lin}} \in \{3, 5, 6, 7, 8, 9, 10, 15, 20, 50\}, \tag{3.16}$$

$$f_r \in \{10, 50, 100, 300, 500\}\,\mathrm{Hz}. \tag{3.17}$$

The resulting 10×5 grid comprised 50 acquisition sessions, each lasting 60 s. The design therefore generated a nominal total of 576 000 input spikes per channel over the complete campaign.

Each condition began with channel calibration using $K = 5$, after which the same regular spike train was applied simultaneously to all 15 inputs. Preliminary tests showed that high firing rates could bias the threshold estimate when spikes were present during calibration. The calibration records were therefore acquired using noise-only signals, allowing the detection performance to be evaluated with thresholds representative of the background noise. This procedure differs from a realistic *in vitro* recording, in which spontaneous spikes may occur during calibration. However, at typical average firing rates of only a few events per second, their influence over a 1 s calibration interval is expected to remain limited. The FPGA-based implementation addresses this issue more explicitly through the spike-masking calibration procedure described in Section 3.5.4.

Within one acquisition, the inter-spike interval (ISI) was constant and determined by

$$\Delta t_{\mathrm{GT}} = t_k - t_{k-1} = \frac{1}{f_r}, \tag{3.18}$$

where t_k is the known generation time of the k-th spike. The 500 Hz condition corresponds to $\Delta t_{\mathrm{GT}} = 2\,\mathrm{ms}$, directly approaching the maximum accepted rate imposed by the firmware refractory interval.

The DAQ and MCU did not share a hardware trigger or absolute time reference. The comparison therefore used the spacing between consecutive captured events rather than their absolute timestamps. For each channel,

$$\Delta \widehat{t}_j = \widehat{t}_j - \widehat{t}_{j-1}, \tag{3.19}$$

where $\widehat{t}_j$ is the timestamp of the j-th captured event. A detected interval was considered consistent with the regular input sequence when

$$\left|\Delta\widehat{t}_j - \Delta t_{\mathrm{GT}}\right| \leq 1\,\mathrm{ms}. \tag{3.20}$$

Intervals satisfying this condition were associated with correctly reconstructed events. Additional detections producing intervals outside the tolerance were classified as false positives, while expected events without a corresponding valid detection were counted as false negatives. Precision, recall, and F1 score were then calculated using Equation 3.15. The MCU timing dispersion was evaluated from the accepted inter-event interval errors,

$$\varepsilon_j = \Delta\widehat{t}_j - \Delta t_{\mathrm{GT}}, \tag{3.21}$$

using their standard deviation. This quantity describes the repeatability of the reconstructed inter-event spacing; it should not be interpreted as an absolute detection latency.

The interval-based procedure removed the unknown constant offset between the stimulus and acquisition clocks and was adequate for the regular trains used in this proof-of-concept campaign. It was not intended for unambiguous one-to-one association under irregular timing or after complex sequences of missed and additional events. The subsequent FPGA campaign therefore adopted an explicit synchronization and event-matching procedure.

3.6.5 FPGA Validation Campaign

The FPGA implementation underwent the principal quantitative validation of the event-driven architecture. It used the common test bench described in Section 3.6.3, extending the MCU proof-of-concept through repeated sessions, a broader experimental design, irregular event timing, amplitude variation after calibration, and direct latency measurements.

The common waveform was distributed to the complete 60-position interface. One channel was assigned to ground, while another channel was excluded because a hardware fault prevented valid event detection. Quantitative event analysis was therefore performed on 58 channels.

Table 3.3 Experimental design of the FPGA validation campaign. Conditions in Studies 1–3 were repeated in five independent sessions.

Study	Purpose	Variable parameters	Fixed conditions
1	Operating-envelope and threshold-multiplier characterization	$K \in \{3,4,5\}$; SNR $\in \{3,4,6,8,10,20\}$; $f_r \in \{5,10,20,50,100,200,300\}$ Hz	$A_{\text{spike}} = 150$ mV; fixed ISI; calibration before every acquisition
2	Robustness to amplitude reduction after calibration	$A_{\text{spike}} \in \{50,75,100,125,150\}$ mV	One calibration per session at $A_{\text{ref}} = 150$ mV; $K = 4$; calibration SNR = 10; $f_r = 50$ Hz
3	Robustness to irregular and clustered timing	Mean $f_r \in \{5,10,20,50,100,200,300\}$ Hz; exponentially distributed ISIs	$K = 4$; SNR = 10; $A_{\text{spike}} = 150$ mV; calibration before every acquisition
4	FPGA-internal and event-driven system latency	Internal latency at one reference condition; system latency for $K \in \{3,4\}$ and SNR $\in \{8,10,20\}$	Diagnostic GPIO exposing channel 33 sticky-bit assertion; $f_r = 50$ Hz; oscilloscope measurements

The campaign comprised four complementary studies, summarized in Table 3.3. Study 1 established the operating envelope and supported selection of the threshold multiplier used in the subsequent protocols. Studies 2 and 3 examined robustness to amplitude variation and irregular timing, respectively. Study 4 evaluated the detection latency of the event-driven FPGA backend.

Study 1: Operating Envelope

Study 1 evaluated how the threshold multiplier, SNR, and firing rate jointly affected detection performance. Its purpose was to identify a suitable sensitivity–noise-rejection trade-off for K and to define the SNR and event-rate region over which the platform operated reliably.

The full $3 \times 6 \times 7$ grid comprised 126 conditions per session and was repeated in five independent sessions, giving 630 acquisitions. The spike amplitude at the motherboard INA input was fixed at 150 mV, and spikes were generated using constant ISIs. A new calibration was performed before every acquisition so that each threshold reflected the noise level of the corresponding stimulus. The host-side refractory filter was disabled during acquisition ($T_{\text{refr}} = 0$) and a 2 ms per-channel rule was applied offline. This separated the response of the physical detector from the effect of the selected post-processing rule. The number of acquisitions sufficiently defines the scale of this factorial study; an aggregate spike count is not required because it varies strongly with firing rate and does not represent an independent number of experimental replicates.

Study 2: Amplitude Reduction after Calibration

Study 2 assessed robustness when spike amplitude decreased after calibration, as may occur because of changes in electrode–cell coupling, baseline conditions, or extracellular signal amplitude.

At the beginning of each session, calibration was performed using $A_{\text{ref}} = 150\,\text{mV}$, SNR = 10, $K = 4$, and $f_r = 50\,\text{Hz}$. The resulting DAC thresholds were held constant while the applied spike amplitude was varied over

$$A_{\text{spike}} \in \{50, 75, 100, 125, 150\}\,\text{mV}. \tag{3.22}$$

Five sessions, each containing all five amplitude conditions, produced 25 acquisitions. Since calibration at the reference condition corresponded to $\sigma_{\text{cal}} = 15\,\text{mV}$, reducing A_{spike} decreased the effective SNR while leaving the threshold unchanged. The protocol therefore represented an intentionally unfavorable loss-of-signal condition rather than a sequence of independently recalibrated operating points.

Study 3: Poisson Spike Timing

Study 3 examined whether the performance obtained with regular spike trains was preserved under irregular timing. ISIs were drawn from an exponential distribution,

$$\Delta t_k \sim \text{Exp}(\lambda = f_r), \qquad \mathbb{E}[\Delta t_k] = \frac{1}{f_r}, \qquad t_k = \sum_{j=1}^{k} \Delta t_j, \tag{3.23}$$

where f_r is the mean firing rate. This timing model generates occasional short intervals and clusters even when the mean rate is moderate, thereby challenging both the 1 ms occupancy window and the post-processing refractory rule.

The mean firing rate was varied over

$$f_r \in \{5, 10, 20, 50, 100, 200, 300\}\,\text{Hz}, \tag{3.24}$$

while $K = 4$, SNR = 10, and $A_{\text{spike}} = 150\,\text{mV}$ were fixed. Each rate was tested in five independent sessions, giving 35 acquisitions, recalibrating before each acquisition.

Study 4: Onboard Latency

Study 4 evaluated the latency of the system to detect a spike and determined whether the temporal response was limited by the FPGA logic or by the preceding analog threshold-detection path. This distinction is relevant when assessing the architecture for real-time and future closed-loop operation. An FPGA general-purpose input/output (GPIO) was configured to expose the assertion of the sticky activity bit for one channel. Two delays were then measured relative to the same GPIO transition:

- FPGA-internal delay τ_{FPGA}, from the signal at the FPGA input pin to sticky-bit GPIO assertion;
- Event-driven system latency τ_{system}, from the positive peak of the DAQ-generated spike to sticky-bit GPIO assertion.

The FPGA-internal delay was measured at $A_{\mathrm{spike}} = 150\,\mathrm{mV}$, SNR = 10, $f_r = 50\,\mathrm{Hz}$, and $K = 4$. Fifteen independently calibrated measurements were acquired using the oscilloscope statistics mode. The event-driven system latency was evaluated for

$$K \in \{3,4\}, \qquad \mathrm{SNR} \in \{8,10,20\}, \tag{3.25}$$

at $f_r = 50\,\mathrm{Hz}$. Ten 10 s sessions were recorded for each of the six conditions. Each session was preceded by a new calibration.

An oscilloscope sampled at $100\,\mathrm{MS\,s^{-1}}$ was used to acquire signals from each measurement point described earlier. The recorded DAQ waveform was processed with a fourth-order zero-phase Butterworth low-pass filter at 3 kHz. Positive peak times (t_{peak}) were identified using a minimum amplitude of 50 mV and a minimum separation of 2 ms. FPGA signal input and GPIO times ($t_{\mathrm{FPGA,in}}$ and t_{GPIO}, respectively) were extracted from low-to-high transitions through a 2.5 V threshold. Each analog peak was paired with the earliest GPIO edge occurring within the following 5 ms, and the event latency was calculated as

$$\tau_{\mathrm{FPGA},j} = t_{\mathrm{GPIO},j} - t_{\mathrm{FPGA,in},j}, \qquad \tau_{\mathrm{system},j} = t_{\mathrm{GPIO},j} - t_{\mathrm{peak},j}. \tag{3.26}$$

3.6.6 FPGA Temporal Alignment and Event Matching

Studies 1–3 required software alignment because the DAQ and FPGA operated from independent oscillators and did not share a hardware start trigger. Every acquisition therefore contained a synchronization preamble followed by the target signal, as shown in Figure 3.7. The preamble used a Barker-13 sequence [48]:

$$\mathbf{b} = [1, 1, 1, 1, 1, 0, 0, 1, 1, 0, 1, 0, 1]. \tag{3.27}$$

Each bit occupied a 15 ms interval. A logical one contained one synchronization spike generated at SNR = 20, whereas a logical zero contained background noise only. The preamble therefore lasted 195 ms and contained nine spikes. It was followed by a 20 ms guard interval and a 60 s target waveform. The first 55 s of the target were retained for analysis; the final 5 s were excluded to avoid end-of-recording and buffer-transition effects.

Because all valid channels received the same waveform, the decoded activity was reduced to a population-level multi-unit-activity sequence by counting the number of simultaneously active channels in each FPGA packet. This population signal provided a more robust alignment reference than any individual channel.

The constant temporal offset was first estimated from the Barker preamble. A candidate alignment was accepted when at least six of its nine spikes matched population events within ±1 ms. It was then validated against the first 60 strict population events in the target waveform, requiring at least 25% agreement within the same tolerance. If this check failed, an exhaustive cross-correlation search over the hardware-constrained interval from −400 ms to −50 ms provided a fallback offset.

Residual oscillator mismatch was corrected through

$$t_{\mathrm{GT}} = \alpha t_{\mathrm{FPGA}} + \beta, \tag{3.28}$$

where β represents the constant offset and α corrects the relative clock rate. The regression used at least 50 matched population events, with a pairing tolerance of

$$\tau_{\mathrm{reg}} = \min(2\,\mathrm{ms}, 0.15 T_{\mathrm{ISI}}). \tag{3.29}$$

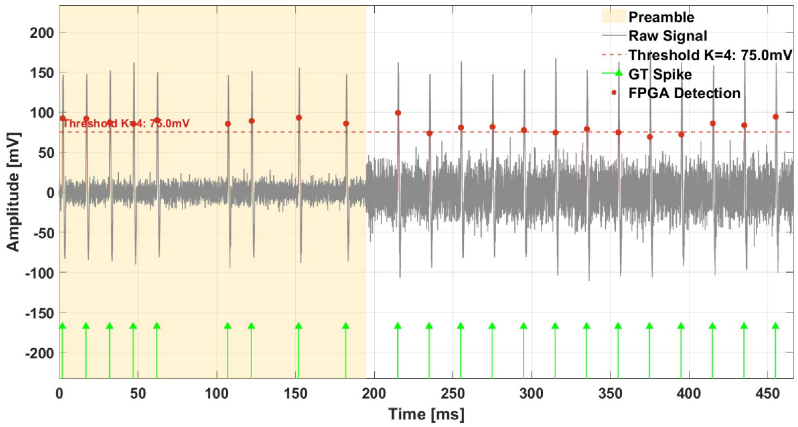


Fig. 3.7 Temporal alignment using a Barker-13 preamble sequence, with each spike representing a 1, while noise represents a 0.

The correction was accepted only when the regression intercept remained within 3 ms of the Barker-derived offset, corresponding to the maximum expected drift of the FPGA clock over the retained analysis interval. The resulting mapping was applied to every per-channel timestamp. For scoring, the nominal ground-truth time was assigned to the positive waveform peak:

$$t_{k,\text{peak}} = t_k + \frac{T}{4} = t_k + 0.5\,\text{ms}. \tag{3.30}$$

Using the peak rather than the spike onset centred the ± 1 ms matching interval over the principal waveform excursion. The resulting TP, FP, and FN counts were converted into the metrics of Equation 3.15.

For channels with at least three matched pairs, jitter was calculated as the standard deviation of the signed temporal errors

$$e_j = t^{\text{corr}}_{\text{det},j} - t_{\text{peak},j}. \tag{3.31}$$

The corresponding mean error was not interpreted as absolute latency because the offset-correction procedure directly affects it. Absolute event-driven system latency was assessed independently through the oscilloscope measurements of Study 4.

3.6.7 Statistical Analysis

For Studies 1–3, the performance metrics were first calculated independently for each of the 58 valid channels. Since all channels received the same waveform and were tested under identical acquisition conditions, they were not treated as independent statistical replicates.

For Study 1, the inferential unit was the mean F_1 score across the 58 channels within each session, yielding five repeated observations for each experimental condition. Inter-session repeatability was summarized using the coefficient of variation

$$CV_{F_1} = 100\,\frac{s_{F_1}}{\overline{F_1}}, \tag{3.32}$$

where $\overline{F_1}$ and s_{F_1} are respectively the mean and sample standard deviation of the five session-level F_1 scores obtained under the same SNR, firing-rate, and threshold-multiplier condition. The coefficient was calculated only for conditions with $\overline{F_1} > 0$. A condition was classified as stable when $CV_{F_1} < 10\%$, according to the predefined engineering repeatability criterion adopted in this study.

For each of the 42 SNR–firing-rate combinations, the three values of K were compared using a Friedman repeated-measures test [49]. This non-parametric procedure was selected because F_1 is bounded between zero and one and only five repeated observations were available per condition. The 42 resulting probabilities formed a single family and were corrected using the Benjamini–Hochberg procedure with a false-discovery rate (FDR) of 0.05 [50]. Kendall's coefficient of concordance W was reported as an effect-size measure for conditions showing a significant threshold-multiplier effect [51]. Studies 2 and 3 were evaluated descriptively.

The FPGA-internal component of Study 4 was summarized descriptively across its 15 measurements. For the event-driven system latency, the event-level latencies obtained within each session were first summarized using their arithmetic mean. The resulting ten session-level means represented the paired inferential observations for each experimental condition.

The effect of SNR was evaluated separately for each value of K using a Friedman repeated-measures test, with Kendall's W reported as an effect-size measure. Post-hoc comparisons among the three SNR levels were performed using exact paired Wilcoxon signed-rank tests [52]. The effect of K was evaluated at each matched

SNR using the same exact paired test. Benjamini–Hochberg FDR correction was applied separately to three predefined families: the two Friedman tests evaluating the SNR effect ($m = 2$), the six post-hoc SNR comparisons ($m = 6$), and the three paired comparisons between $K = 3$ and $K = 4$ ($m = 3$).

3.7 Results

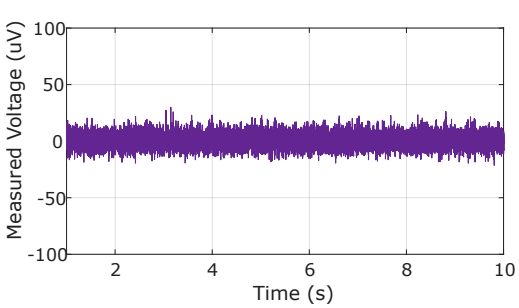
3.7.1 Preliminary Single-Channel AFE Results

For the standalone electrical and biological AFE tests, the amplifiers were configured with a total nominal gain of approximately 2000. The electrical tests confirmed that the AFE could acquire artificial signals at the intended microvolt-level input scale. As shown in Figure 3.8a, under the zero-input condition, the noise referred to the AFE input was approximately 5.1 µV RMS.

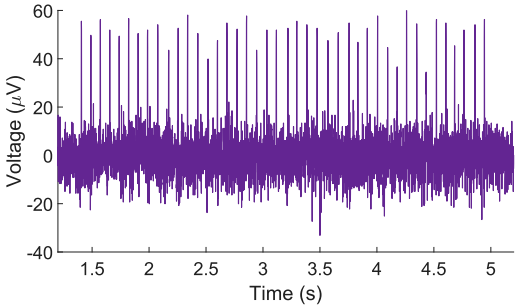
Figure 3.8b displays the recordings after the injection of a pulse train, with 50 µV, 1 ms pulses. The measured pulse peaks had a mean amplitude of 50.69 µV and a standard deviation of 6.52 µV, showing that the applied signal amplitude was recovered with the correct order of magnitude.

During *in vitro* experiments, extracellular spike-like activity was observed when the AFE was connected to four selected electrodes of the neuronal culture. Across all recorded channels, displayed in Figure 3.8c, the input-referred RMS noise remained below 10 µV. Unlike the zero-input laboratory measurement, this value includes not only the contribution of the electronic chain, but also noise associated with the MEA electrode, the electrode–electrolyte interface, and the biological preparation.

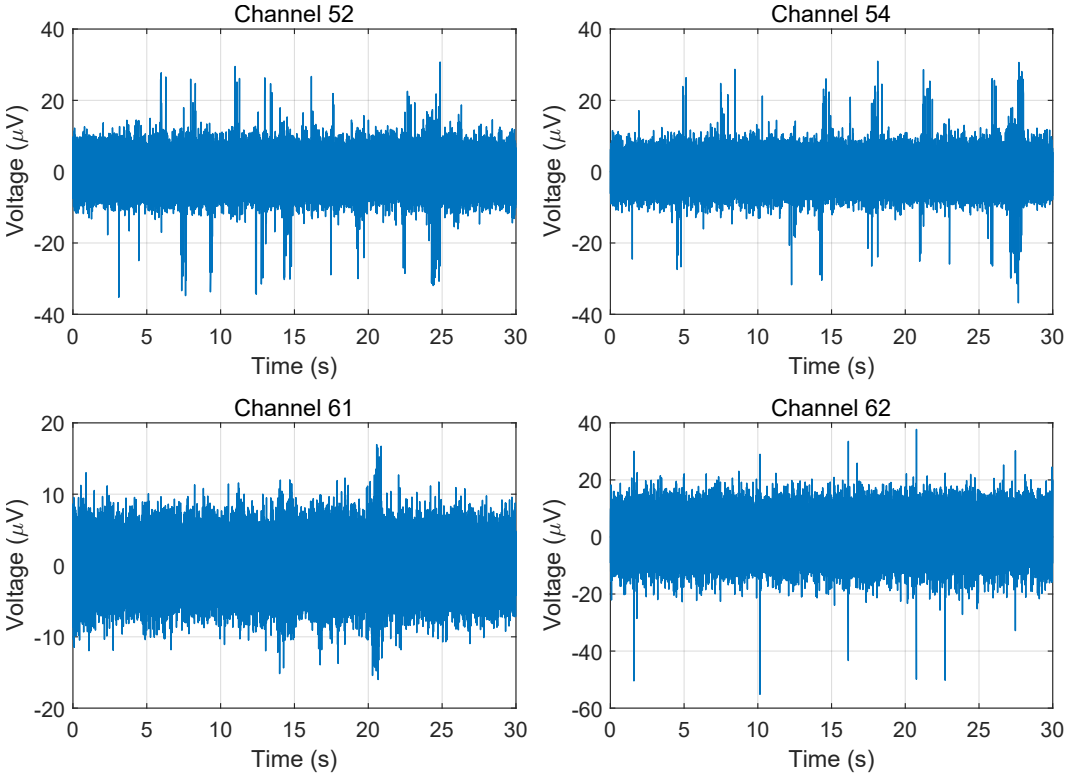
Neural spike-like events were distinguishable from the background in all available active-channel recordings. Across the four analysed electrodes, the mean negative-peak amplitude ranged from approximately 22.3 µV to 37.4 µV in magnitude, while the estimated mean event rate ranged from 1.0 Hz to 10.6 Hz. The recordings also contained intervals of visibly increased event density. These patterns are consistent with burst-like activity, although no formal burst-detection algorithm was applied and they should therefore not be interpreted as a quantitative characterization of network bursting.



(a) Zero-input recording.



(b) 50 µV, 1 ms artificial-pulse recording.



(c) 30 s input-referred extracellular recordings from four selected MEA electrodes.

Fig. 3.8 Preliminary characterization of the single-channel analog front end. (a) Zero-input electrical recording, showing the background noise of the electronic acquisition chain. (b) Response to controlled microvolt-scale artificial pulse train. (c) Representative *in vitro* recordings obtained from four selected electrodes of a passive MEA.

Taken together, these measurements show that the custom analog path could resolve extracellular signals within the expected amplitude and firing-rate ranges while maintaining a noise level below 10 µV RMS under both electrical and biological test conditions. This represents a favorable result for a preliminary single-channel

demonstrator, demonstrating compatibility with recordings in which *in vitro* neural extracellular spikes typically have amplitudes of the order of tens of µV.

The evidence nevertheless remains preliminary because it was obtained from a limited number of electrodes and biological preparations, without a simultaneous reference acquisition or formal biological event-identification protocol. The results therefore confirmed the low-noise properties of the single-channel AFE and established feasibility of neural recordings, but did not quantify absolute accuracy, inter-culture repeatability, or long-term stability.

3.7.2 MCU Proof-of-Concept Results

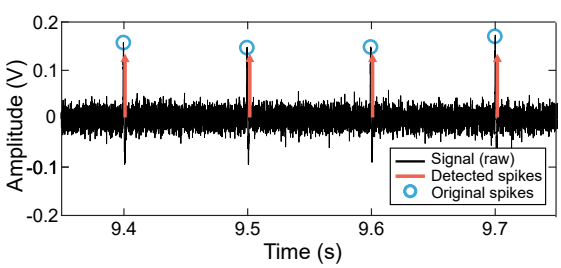
The MCU-based platform, tested according to the protocol presented in Section 3.6.3, successfully reconstructed the regular synthetic spike trains applied simultaneously to all 15 tested channels.

Figure 3.9 provides a representative overview of the MCU results. The offline overlay in Figure 3.9a qualitatively shows the correspondence between the generated spikes and the reconstructed events for a 400 ms interval. Figure 3.9b instead shows an example of a raster plot reconstructed by the host interface.

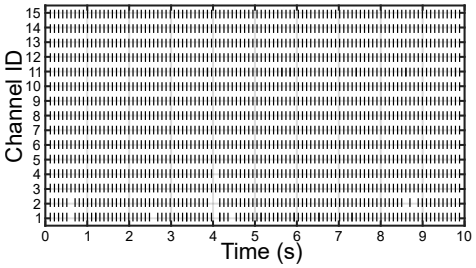
For each acquisition, detection performance was obtained averaging the metrics between all 15 channels. Performance showed a strong dependence on signal quality. At $f_r = 10\,\text{Hz}$, the F_1 score increased sharply as the linear SNR increased, as shown in Figure 3.9c. It exceeded approximately 0.90 from $\text{SNR}_{\text{lin}} \simeq 8$, reached approximately 0.98 at $\text{SNR}_{\text{lin}} = 10$, and approached unity under the highest-SNR conditions. With lower SNR, the spike amplitude was close to the threshold and noise fluctuations and channel-level variations could prevent some spikes from crossing the detector boundary, decreasing performance as expected. Once the spike amplitude was

Table 3.4 Representative MCU proof-of-concept results at $\text{SNR}_{\text{lin}} = 10$, aggregated over firing rates up to 300 Hz.

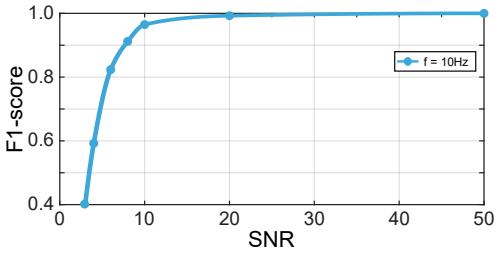
Metric	**Result**
Recall	0.96
Precision	≈ 1.00
F_1 score	0.98
Mean inter-event interval-error standard deviation	0.66 ms



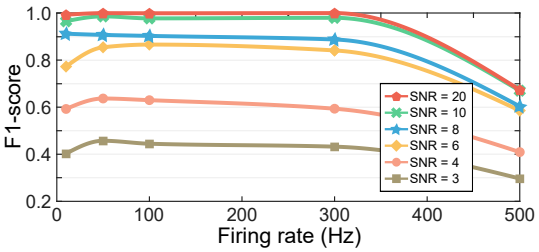
(a) Offline comparison of generated and detected events.



(b) 15-channel raster reconstruction.



(c) F_1 score as a function of linear SNR.



(d) F_1 score as a function of firing rate.

Fig. 3.9 Results of the MCU-based 15-channel proof-of-concept validation. (a) Offline overlay of a representative 400 ms synthetic input segment at $\mathrm{SNR}_{\mathrm{lin}} = 10$ and $f_r = 10\,\mathrm{Hz}$, showing the noisy input waveform in black, the known spike occurrences as blue circles, and the reconstructed events as red arrows. The overlay was generated during post-processing for qualitative visualization and does not constitute an absolute temporal-alignment measurement. (b) Representative host-side raster plot obtained during a 10 s acquisition with all 15 channels stimulated concurrently. (c) Channel-averaged F_1 score as a function of the dimensionless linear amplitude ratio $\mathrm{SNR}_{\mathrm{lin}}$, evaluated at $f_r = 10\,\mathrm{Hz}$. (d) Channel-averaged F_1 score as a function of firing rate for selected values of $\mathrm{SNR}_{\mathrm{lin}}$. For representation clarity, not all the acquisition sessions are shown.

sufficiently above this boundary, detection became substantially insensitive to further increases in SNR.

The conservative $K = 5$ threshold and the 2 ms refractory interval strongly suppressed FPs. The false-positive rate remained below approximately 0.05 events s^{-1} $\mathrm{channel}^{-1}$, even under the lowest-SNR condition. Precision was consequently close to unity over most of the tested operating range. The reduction in F_1 observed at low SNR was therefore caused mainly by false negatives rather than by noise-induced false positives. In practical terms, the adopted configuration behaved as a conservative detector: unreliable signals were more likely to produce missed events than spurious activity.

The firing-rate sweep, reported in Figure 3.9d, showed that performance remained substantially stable up to 300 Hz. At $\mathrm{SNR}_{\mathrm{lin}} = 10$, aggregation over firing rates up to 300 Hz yielded approximately 96% recall, near-unity precision, and an F_1 score of 0.98. As expected, a marked performance reduction occurred at $f_r = 500\,\mathrm{Hz}$. Under this condition, the nominal ISI was

$$\Delta t_{\mathrm{GT}} = \frac{1}{f_r} = 2\,\mathrm{ms}, \tag{3.33}$$

which was equal to the refractory interval enforced by the firmware. Small variations in threshold-crossing time could therefore cause a subsequent event to occur before the channel had been re-enabled, producing a false negative. The decrease at 500 Hz consequently identifies the boundary imposed by the selected refractory policy.

Temporal consistency was evaluated from the standard deviation of the accepted inter-event interval errors defined in Section 3.6.4. The timing dispersion remained below 1 ms throughout the principal operating region. Representative values were approximately 0.59 ms at 50 Hz and 0.68 ms at 300 Hz, with an average value of approximately 0.66 ms over firing rates up to 300 Hz. These values describe the repeatability of the reconstructed inter-event spacing and include the effect of the 1 ms transmitted activity bins.

The output data-rate reduction obtained for the MCU-based platform is reported separately in Section 3.7.4.

3.7.3 FPGA-Based Validation Results

The FPGA-based implementation underwent the principal quantitative characterization of the event-driven architecture. The four studies introduced in Section 3.6.5 evaluated complementary aspects of its operation: the detection envelope and threshold selection, robustness to signal-amplitude variation after calibration, performance under irregular spike timing, and onboard detection latency. The aggregation procedures and statistical units are those defined in Section 3.6.7. Studies 1 and 4 included inferential analysis, whereas Studies 2 and 3 were evaluated descriptively.

Table 3.5 summarizes the main findings. The results refer to the controlled post-AFE validation path and should not be interpreted as an end-to-end characterization from microvolt-level biological signals.

Study 1: Operating Envelope and Threshold Selection

Study 1 characterized the detection operating envelope as a function of the linear SNR, firing rate, and threshold multiplier K, according to the protocol defined in Section 3.6.5. Detection performance was determined primarily by the relationship between the threshold and the applied spike amplitude. Since the threshold is proportional to $K\widehat{\sigma}_n$, this relationship can be approximated as

$$\frac{T}{A_{\text{spike}}} \simeq \frac{K}{\text{SNR}_{\text{lin}}}. \tag{3.34}$$

Increasing K therefore improves rejection of false positives but shifts the detection boundary toward higher SNR_{lin}.

For every tested K, no events were detected at $\text{SNR}_{\text{lin}} = 3$ and 4. For firing rates between 5 Hz and 100 Hz, $K = 3$ yielded $F_1 \simeq 0.92$–0.93 already at $\text{SNR}_{\text{lin}} = 6$. With $K = 4$, performance remained poor at this SNR but exceeded approximately 0.94 from $\text{SNR}_{\text{lin}} = 8$. With $K = 5$, no events were detected at $\text{SNR}_{\text{lin}} = 6$; F_1 ranged from approximately 0.40 to 0.59 at $\text{SNR}_{\text{lin}} = 8$, and reached values near 0.97 from $\text{SNR}_{\text{lin}} = 10$. The resulting operating envelopes are shown in Figure 3.10.

Within the high-performance region, firing rate had a limited influence up to 100 Hz. At 200 Hz, high F_1 values were still obtained when the SNR was sufficiently large. At 300 Hz, however, F_1 was zero or close to zero for all three threshold multipliers up to $\text{SNR}_{\text{lin}} = 10$. Detection recovered at $\text{SNR}_{\text{lin}} = 20$, reaching $F_1 = 0.947$, 0.946, and 0.754 for $K = 3$, 4, and 5, respectively.

Table 3.5 Summary of the FPGA-based validation results.

Study	Main result	Key outcome
1	For $f_r \leq 100$ Hz, reliable detection was obtained from $\text{SNR}_{\text{lin}} = 6$, 8, and 10 for $K = 3$, 4, and 5, respectively.	$K = 3$ provided the widest operating envelope; dense regular activity degraded calibration at the highest firing rates.
2	F_1 decreased to 0.570 ± 0.302 at $A_{\text{spike}} = 50$ mV, but exceeded 0.93 from 75 mV.	Amplitude reduction primarily increased false negatives; precision remained high.
3	Under Poisson timing, F_1 decreased from 0.935 ± 0.117 at 5 Hz to 0.639 ± 0.130 at 300 Hz.	Closely spaced events were merged by the occupancy window or removed by refractory filtering.
4	$\tau_{\text{FPGA}} = (80.41 \pm 3.39)$ ns. Across sessions, mean τ_{system} ranged from approximately 0.34 ms to 1.45 ms.	Event-driven system latency was dominated by threshold-detection path rather than by the FPGA logic.

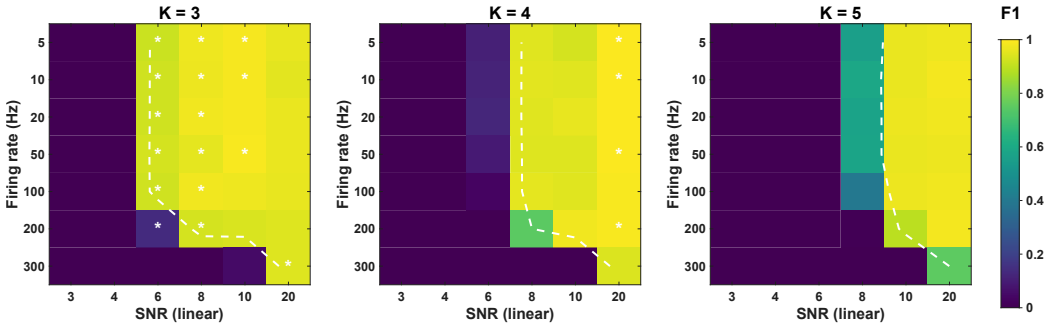


Fig. 3.10 Detection operating envelope obtained in FPGA Study 1. The heatmaps report the grand-average F_1 score across the tested $\mathrm{SNR}_{\mathrm{lin}}-f_r$ space for each threshold multiplier K. The white dashed contour identifies the adopted $F_1 \geq 0.75$ operating threshold. Asterisks mark conditions in which the corresponding K value produced the statistically highest session-level F_1 score after FDR correction.

Inspection of the calibration traces indicated that the high-rate degradation originated primarily from the limited amount of uncontaminated baseline available during calibration, despite the spike-masking stage. At $f_r = 300\,\mathrm{Hz}$, the regular ISI was only approximately 3.33 ms. Since samples exceeding the preliminary $4\widehat{\sigma}_{\mathrm{perc}}$ boundary were dilated by ± 15 samples, the exclusion regions associated with consecutive spikes covered most of the 100 ms calibration record. The number of surviving samples could therefore approach or fall below the minimum required for the second-stage MAD estimate, causing the algorithm to rely on the preliminary percentile-based estimate. Moreover, sub-threshold waveform tails could remain outside the exclusion mask and bias this estimate upward. The resulting DAC threshold could consequently approach or exceed the spike excursion. The high-rate failure should therefore not be interpreted as saturation of the parallel FPGA event-capture logic.

At low SNR, the decrease in F_1 was generally dominated by reduced recall, while precision remained comparatively high. The detector therefore behaved conservatively near the operating boundary, preferentially missing low-amplitude events rather than generating large numbers of false detections. Under the most adverse combinations of high firing rate and insufficient SNR, both precision and recall deteriorated. Complete condition-level F_1, precision, and recall values are reported in Tables A.1–A.3. Alternative representations of the same dataset as a function of SNR and firing rate are provided in Figure A.1.

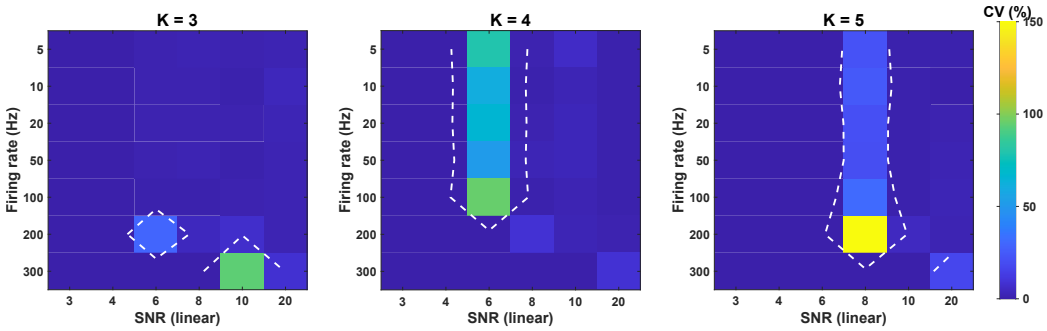


Fig. 3.11 Inter-session repeatability in FPGA Study 1, expressed as the CV of the session-level F_1 score. The white dashed contour marks the predefined $CV = 10\%$ threshold.

Repeatability across the five sessions could be quantified for 69 of the 126 experimental conditions. The CV was undefined for the remaining conditions because all five sessions produced $F_1 = 0$. Among the computable conditions, 55 of 69 (79.7%) satisfied the predefined $CV < 10\%$ criterion, and the median CV was 1.6%. As shown in Figure 3.11, higher variability was concentrated near the SNR-dependent detection boundaries and at the highest firing rates, where small differences in event detection produced large relative changes in F_1.

The effect of K was statistically significant after FDR correction in 20 of the 42 SNR–firing-rate combinations. Among these conditions, $K = 3$ produced the highest session-level F_1 in 16 cases and $K = 4$ in four cases; $K = 5$ was never statistically preferred. Kendall's W ranged from 0.760 to 1.000, with a mean of 0.900, indicating a large and consistent threshold-multiplier effect in the significant conditions. The full list of significant comparisons, corrected p-values, and effect sizes is reported in Table A.5. Conditions exceeding the repeatability threshold are listed separately in Table A.4.

For conditions with effective detection and $f_r \leq 100\,\text{Hz}$, temporal jitter remained below 0.4 ms, as shown in Figure 3.12. Across the tested SNR values, jitter ranged from 0.175 ms to 0.367 ms for $K = 3$, from 0.168 ms to 0.320 ms for $K = 4$, and from 0.143 ms to 0.237 ms for $K = 5$. Higher thresholds produced lower jitter in some conditions close to the detection boundary, but the trend was not monotonic across the complete SNR range and was accompanied by reduced sensitivity. The corresponding numerical values are reported in Table A.6.

Overall, $K = 3$ provided the widest low-SNR operating envelope, whereas increasing K progressively traded sensitivity for more conservative thresholding.

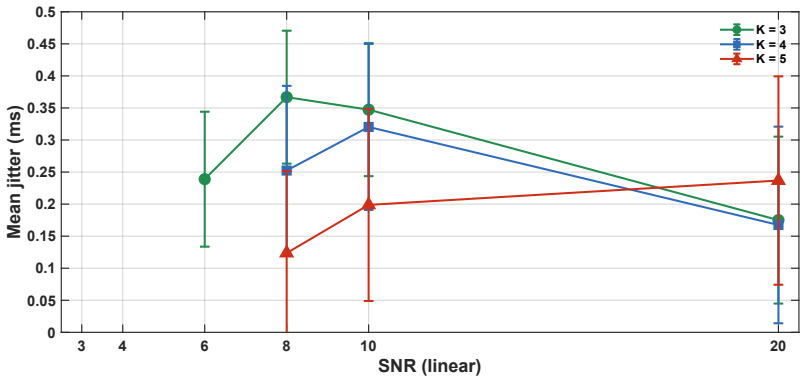


Fig. 3.12 Mean temporal jitter in FPGA Study 1 as a function of $\mathrm{SNR}_{\mathrm{lin}}$ and threshold multiplier K. The analysis only included firing rates up to 100 Hz and to conditions with effective detection. Error bars indicate the standard deviation.

The $K = 4$, $\mathrm{SNR}_{\mathrm{lin}} = 10$, and $f_r = 50\,\mathrm{Hz}$ condition used in Studies 2 and 3 was a preselected intermediate operating point within the stable region.

Study 2: Amplitude Reduction after Calibration

Study 2 evaluated the response to a reduction in spike amplitude after the detection thresholds had already been calibrated, according to the protocol defined in Section 3.6.5. Because the same calibration-derived noise estimate and DAC thresholds were retained throughout the amplitude sweep, the experiment isolated robustness to post-calibration signal variation.

Detection performance decreased primarily through a reduction in recall. At $A_{\text{spike}} = 50\,\text{mV}$, corresponding to an effective $\text{SNR}_{\text{lin}} = 3.33$, the F_1 score was 0.570 ± 0.302. Precision remained high at 0.974 ± 0.119, whereas recall decreased to 0.461 ± 0.294. The fixed thresholds therefore continued to reject background fluctuations, but a substantial fraction of the reduced-amplitude spikes no longer crossed the detection boundary.

Performance recovered markedly at $A_{\text{spike}} = 75\,\text{mV}$, where $F_1 = 0.938 \pm 0.131$. Between 100 mV and the 150 mV calibration amplitude, F_1 remained between approximately 0.95 and 0.97, with closely balanced precision and recall. The corresponding trend is shown in Figure 3.13.

The partial detection obtained at $A_{\text{spike}} = 50\,\text{mV}$ is not directly comparable with the low-SNR conditions of Study 1. In Study 2, the threshold had been calibrated previously at $\text{SNR}_{\text{lin}} = 10$ and was then kept fixed; when the amplitude was reduced, the spike peak remained close enough to this threshold for a fraction of events to be detected. The effective $\text{SNR}_{\text{lin}} = 3.33$ therefore describes the reduced signal relative to the calibration noise, rather than a new calibration performed at that SNR.

Temporal jitter ranged from 0.232 ms to 0.396 ms and did not vary monotonically with spike amplitude. Amplitude reduction therefore affected the probability of detecting an event more clearly than the temporal dispersion of the events that remained detectable. Complete F_1, precision, recall, and jitter values are reported in Table A.7.

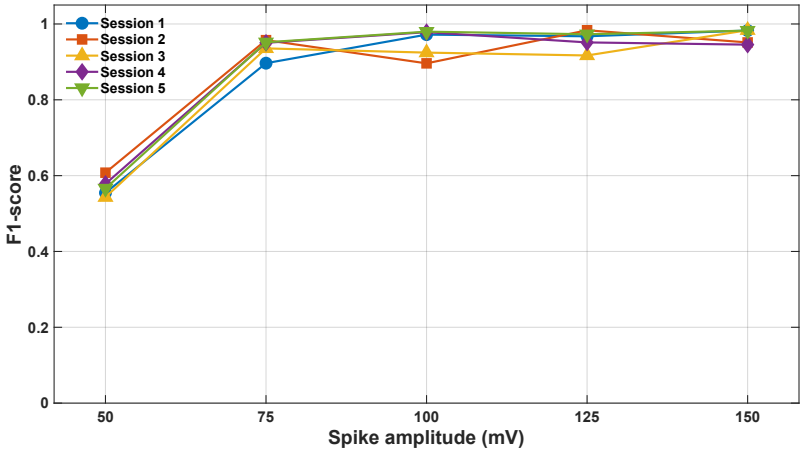


Fig. 3.13 Detection performance in FPGA Study 2 as a function of spike amplitude after a single calibration at $A_{\text{ref}} = 150\,\text{mV}$. Each curve reports the channel-averaged F_1 score for one experimental session. The calibration-derived noise estimate and DAC thresholds were held fixed throughout the amplitude sweep.

Study 3: Poisson-Distributed Spike Timing

Study 3 evaluated detection performance under the Poisson-distributed spike timing defined in Section 3.6.5. In contrast to the regular trains used in Study 1, the sequences contained variable ISIs and therefore reproduced event clustering more realistically.

The F_1 score decreased progressively as the mean firing rate increased, as shown in Figure 3.14. It was 0.935 ± 0.117 at 5 Hz, remained above 0.90 up to 50 Hz, and decreased to 0.855 ± 0.075 at 100 Hz. At 200 Hz and 300 Hz, F_1 reached 0.779 ± 0.069 and 0.639 ± 0.130, respectively.

The degradation was predominantly recall-limited. Recall decreased from 0.930 ± 0.120 at 5 Hz to 0.497 ± 0.111 at 300 Hz, whereas precision remained between approximately 0.92 and 0.98 throughout the tested range. Increasing event density therefore caused expected spikes to be omitted more frequently than additional events to be spuriously generated.

For a Poisson process with mean firing rate f_r, the probability of an ISI shorter than a duration T is

$$P(\Delta t < T) = 1 - \exp(-f_r T)\,. \tag{3.35}$$

At $f_r = 300\,\text{Hz}$, approximately 26% of the intervals are expected to be shorter than the 1 ms occupancy window, and approximately 45% shorter than the 2 ms offline

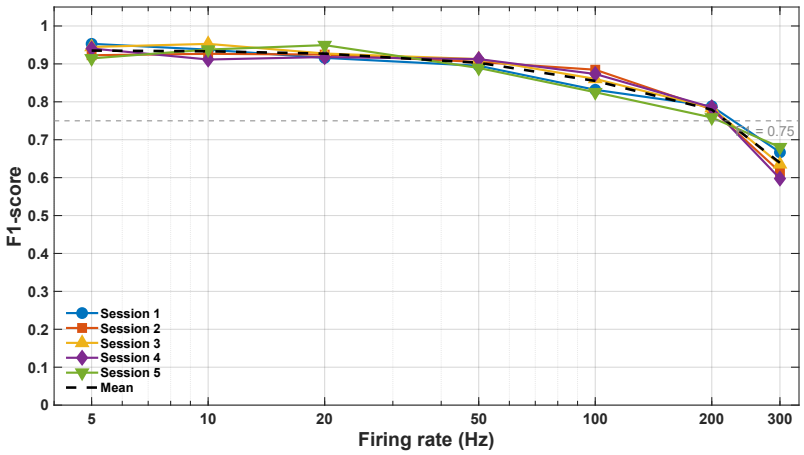


Fig. 3.14 Detection performance in FPGA Study 3 as a function of mean firing rate under Poisson-distributed spike timing. Colored curves report the channel-averaged F_1 score for the individual sessions, while the dashed black curve indicates the session average. The progressive reduction at high mean firing rates is primarily associated with closely spaced events affected by the occupancy-window and refractory constraints.

refractory interval. Multiple closely spaced spikes can consequently be represented by a single occupied window or reduced to one event during post-processing. This mechanism explains the progressive loss of recall and should not be interpreted as saturation of the parallel FPGA event-capture logic.

Temporal jitter remained comparatively stable, ranging from 0.344 ms to 0.432 ms. The principal effect of increasing the mean firing rate was therefore a loss of reconstructed event multiplicity rather than a progressive deterioration in the timing of matched events. Complete numerical results are reported in Table A.8.

Study 4: Onboard Detection Latency

Study 4 separated the latency introduced by the registered FPGA processing path from the latency of the complete event-driven detection chain, according to the protocol described in Section 3.6.5. The event-driven system latency was measured from the positive peak of the synthetic waveform applied at the motherboard input to the assertion of the corresponding FPGA activity bit.

The measured FPGA-internal latency agreed with the interval predicted from the registered conditioning pipeline. At the 50 MHz system clock, the total theoretical delay was within 80 ns to 120 ns, as described in Section 3.5.3. Across 15 sessions,

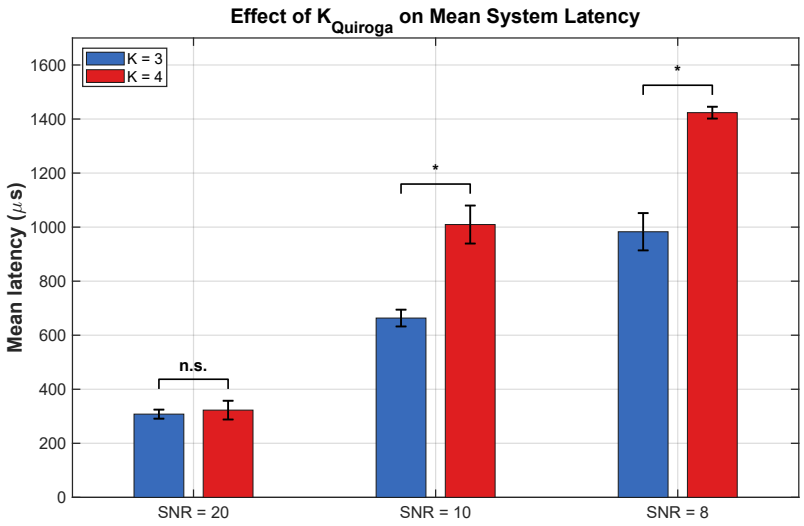


Fig. 3.15 Event-driven system latency in FPGA Study 4 as a function of $\mathrm{SNR}_{\mathrm{lin}}$ and threshold multiplier K. Bars report the mean of the ten session-level mean latencies, and error bars indicate their standard deviation across sessions. Brackets compare $K = 3$ and $K = 4$ within each SNR condition using exact Wilcoxon signed-rank tests after FDR correction; an asterisk denotes a significant difference and n.s. a non-significant difference. The effect of SNR was evaluated separately within each K condition and is reported in the text and in Appendix A.1.3. The $\mathrm{SNR}_{\mathrm{lin}} = 20$ measurements involved a different comparator-output transition and therefore require separate physical interpretation.

the mean latency was

$$\tau_{\mathrm{FPGA}} = (80.41 \pm 3.39)\,\mathrm{ns}, \tag{3.36}$$

where the uncertainty denotes the standard deviation across the session means. The within-session standard deviation remained below 20 ns in every session. The measured latency was therefore consistent with the implemented digital pipeline and constituted only a small fraction of the complete event-driven system response. Complete session-level results are reported in Table A.9.

For the event-driven system latency τ_{system}, the event-level measurements obtained within each session were summarized by their arithmetic mean. The ten resulting session-level means were then aggregated as mean and standard deviation across sessions.

For $K = 3$, the event-driven system latency was $(343.2 \pm 63.0)\,\mu\mathrm{s}$ at $\mathrm{SNR}_{\mathrm{lin}} = 20$, $(701.4 \pm 46.7)\,\mu\mathrm{s}$ at $\mathrm{SNR}_{\mathrm{lin}} = 10$, and $(1005.2 \pm 95.6)\,\mu\mathrm{s}$ at $\mathrm{SNR}_{\mathrm{lin}} = 8$. For $K = 4$, the corresponding values were $(347.1 \pm 47.2)\,\mu\mathrm{s}$, $(1042.6 \pm 79.2)\,\mu\mathrm{s}$, and $(1445.2 \pm 34.7)\,\mu\mathrm{s}$, respectively. The results are summarized in Figure 3.15.

The effect of SNR on the session-level mean latency was significant for both threshold multipliers ($\chi^2_{\text{Friedman}} = 20$, $p < 0.0001$, $W = 1.000$). All three pairwise SNR comparisons within each value of K remained significant after FDR correction ($p_{\text{FDR}} = 0.002$, $|r_{\text{rb}}| = 1.000$). The ordering of the observed event-driven system responses was therefore consistent across all ten sessions.

Threshold selection also affected the mean latency at the lower SNR values. At $\text{SNR}_{\text{lin}} = 8$ and 10, $K = 3$ produced a significantly lower session-level mean latency than $K = 4$ ($p_{\text{FDR}} = 0.003$, $|r_{\text{rb}}| = 1.000$ in both cases). At $\text{SNR}_{\text{lin}} = 20$, the difference was not significant ($p_{\text{FDR}} = 1.000$, $|r_{\text{rb}}| = 0.000$).

Since τ_{FPGA} remained effectively constant, the SNR- and K-dependent variations originated upstream of the registered FPGA logic and mainly reflected the time required for the finite-slope waveform to cross the programmed analog threshold. Complete per-session event-driven system-latency measurements are reported in Tables A.10 and A.11. The complete inferential results are provided in Tables A.12–A.14.

3.7.4 Output Data Rate

One of the main system-level results of the event-driven architecture was the reduction in output payload relative to continuous multichannel sampling. For a continuously sampled system, the ideal packed-data rate was calculated as

$$R_{\text{cont}} = \frac{N_{\text{ch}} f_s b}{8}, \tag{3.37}$$

where N_{ch} is the number of channels, f_s the sampling frequency, and $b = 12$ the ADC resolution. This calculation assumes ideal 12-bit packing and excludes communication-protocol overhead, thereby providing a lower-bound estimate of the continuous-stream payload.

The MCU implementation transmitted a 1003-byte packet every 500 ms, corresponding to $2.006\,\text{kB}\,\text{s}^{-1}$. At the experimentally adopted sampling rate of $10\,\text{kS}\,\text{s}^{-1}$, continuous acquisition of the same 15 channels would generate $225\,\text{kB}\,\text{s}^{-1}$, corresponding to a payload reduction of approximately 99.1%.

The FPGA implementation generated one 12-byte record per 1 ms occupancy window, resulting in a fixed output rate of $12\,\text{kB}\,\text{s}^{-1}$. Continuous acquisition of 60

Table 3.6 Comparison between the event-driven output rates and ideal continuous 12-bit acquisition. The continuous rates assume ideal sample packing and exclude protocol overhead. The $20\,\mathrm{kS\,s^{-1}}$ and $50\,\mathrm{kS\,s^{-1}}$ cases are analytical extrapolations.

Platform	**Event output** ($\mathrm{kB\,s^{-1}}$)	f_s ($\mathrm{kS\,s^{-1}}$)	**Continuous rate** ($\mathrm{kB\,s^{-1}}$)	**Reduction**
MCU, 15 channels	2.006	10	225	99.1%
MCU, 15 channels	2.006	20	450	99.6%
MCU, 15 channels	2.006	50	1125	99.8%
FPGA, 60 channels	12.0	10	900	98.7%
FPGA, 60 channels	12.0	20	1800	99.3%
FPGA, 60 channels	12.0	50	4500	99.7%

channels at $10\,\mathrm{kS\,s^{-1}}$ would instead require $900\,\mathrm{kB\,s^{-1}}$, corresponding to a reduction of approximately 98.7%. The packet formats and corresponding comparisons are summarized in Table 3.6.

For both implementations, the output rate was determined by the fixed record format and time-window duration rather than by the number of detected events. Consequently, multiple channels could be active within the same time window without increasing the transmitted payload. The higher-rate continuous sampling cases in Table 3.6 illustrate the linear scaling of the conventional payload with f_s; they do not represent separate experimental operating conditions.

3.7.5 Prototype Cost and Accessibility

Limiting hardware cost was an explicit design objective of the platform. The estimated hardware cost of the final 60-channel prototype was approximately USD 3,000, based on the cost of 60 assembled analog front-end boards, the motherboard, both MCU and FPGA development boards, and associated accessories, including adapter PCBs, cables, and additional components. This estimate excludes development time, personnel costs, and earlier hardware revisions. It reflects low-volume prototype assembly rather than an optimized production bill of materials [6].

For context, previous studies reported total costs of approximately USD 30,000–50,000 for commercial 60-channel MEA instrumentation, with acquisition hardware alone exceeding USD 35,000 in representative configurations [26, 27]. More recent work similarly described commercial *in vitro* MEA platforms as costing several tens

of thousands of euros [28]. These figures are historical, configuration-dependent estimates rather than current commercial quotations.

The prototype expenditure and commercial purchase prices have different cost bases. Commercial prices also reflect product development and qualification, manufacturing and quality assurance, software maintenance, distribution, warranty, technical support, business overheads, and profit margins. Commercial systems generally provide continuous high-resolution waveform recording, stimulation capabilities, mature analysis software, and validated experimental workflows, whereas the present prototype focuses on hardware-level event extraction and spatiotemporal monitoring.

The reported estimate therefore documents the hardware expenditure required to assemble this research prototype. Assessing whether the platform reduces the financial barrier to MEA experiments would require accounting for its assembly, integration, calibration, validation, and maintenance costs, alongside the functionality required by the intended application.

3.8 Discussion

The staged validation addressed complementary parts of this architecture. The preliminary AFE experiments showed that the custom analog chain could resolve microvolt-scale artificial signals and extracellular spike-like activity with input-referred RMS noise below $10\,\mu V$. The MCU implementation then demonstrated the complete 15-channel event-driven workflow, including calibration, hardware event capture, packetization, host-side reconstruction, and raster visualization. Finally, the FPGA backend extended the concept to 60 concurrently monitored channels and enabled a more systematic investigation of operating limits, parameter selection, temporal statistics, and propagation latency.

The transition from MCU to FPGA was primarily motivated by scalability. Although the MCU captured timestamps through hardware peripherals, channel management and buffer operations remained partly sequential and dependent on the available timer resources. The FPGA instead processed all event lines in parallel against a common time base. Simultaneous activity on several electrodes was therefore retained without channel prioritization, while the packet rate remained

independent of the number of active channels within each window. This behaviour is particularly advantageous during population bursts, when conventional address-by-address processing may experience its highest instantaneous traffic.

The FPGA campaign was intentionally broader than a nominal performance test. Its purpose was to identify how the platform should be configured, where its current boundaries arise, and which modifications would most effectively extend its operating range. Detection was governed by the relationship among spike amplitude, background noise, and threshold multiplier. Lower values of K increased sensitivity, whereas higher values produced more conservative detection and required greater signal separation. In the tested conditions, $K = 3$ provided the widest operating envelope, while $K = 4$ offered an intermediate operating point with high performance at $\mathrm{SNR}_{\mathrm{lin}} = 10$.

The experiments also distinguished limitations of calibration from limitations of event encoding. Dense regular activity could contaminate the short calibration record even after spike masking, increasing the estimated noise and the programmed threshold. Changes occurring after calibration produced a complementary effect: when spike amplitude decreased while the threshold remained fixed, performance became predominantly recall-limited. Both results indicate that long-duration recordings would benefit from periodic or adaptive recalibration based on baseline noise, signal amplitude, electrode impedance, event-rate changes, or the number of uncontaminated samples available during calibration.

Under irregular timing, the main limitation arose from the 1 ms occupancy representation. Events occurring on different channels within the same window were retained independently, whereas multiple events from the same channel were represented by a single asserted bit. Closely spaced spikes could also be removed by the optional refractory rule. The resulting loss of recall at high Poisson rates therefore reflected the selected compression policy rather than saturation of the FPGA logic. Shorter windows, local event counters, or individual event timestamps could extend the usable rate range at the cost of increased output traffic.

The latency measurements similarly separated an architectural strength from an implementation limitation. The FPGA contributed approximately 80 ns, confirming that the event capture introduced a negligible delay. Across the ten sessions, the mean event-driven system latency ranged from approximately 0.34 ms to 1.45 ms and increased as SNR decreased and the threshold multiplier increased. The overall

response was dominated by analog threshold crossing and by the low-power Schmitt-trigger stage rather than by the FPGA logic. Replacing this stage with a faster dedicated comparator and an optimized hysteresis network could therefore reduce the event-driven system latency substantially without changing the digital architecture.

The most significant system-level outcome was the output-data reduction. At $10\,\mathrm{kS\,s^{-1}}$, the MCU and FPGA formats reduced the ideal continuous 12-bit payload by approximately 99.1% and 98.7%, respectively. Moreover, the fixed time–channel mask bounded the communication rate even during synchronous multichannel activity. This simplifies buffering, communication scheduling, storage, and real-time processing, and makes the event-driven architecture particularly attractive when the biological endpoint is firing activity, burst timing, recruitment, propagation, or network synchrony.

Limiting hardware expenditure was a complementary design objective. The complete 60-channel prototype, including both digital backends, had an estimated hardware cost of approximately USD 3,000. As discussed in Section 3.7.5, this estimate reflects low-volume prototype assembly and excludes development time, personnel costs, and the costs of earlier hardware revisions. The overall cost of adopting the platform would also depend on integration, calibration, validation, and maintenance, alongside the functionality required by the intended application.

For these applications, the proposed platform may be more appropriate than a conventional continuous acquisition system rather than merely a lower-cost substitute. Processing events directly avoids transferring and analysing large volumes of baseline samples that are subsequently discarded. The resulting data structure is immediately compatible with raster analysis, network-state estimation, and event-triggered control. These advantages become increasingly consequential as recording duration and channel count increase, particularly in high-density MEA systems, where the absolute bandwidth and storage savings can become substantial.

The corresponding drawback is the intentional loss of waveform information. The event stream does not preserve spike morphology, continuous amplitude, or the possibility of retrospective filtering, re-thresholding, and waveform-based spike sorting. The platform is therefore not suitable when unit discrimination, signal morphology, artifact inspection, or complete signal reconstruction are primary objectives. This is an intrinsic trade-off of the event-driven representation rather than a defect that can be removed without increasing the transmitted information.

A hybrid operating mode could nevertheless mitigate this limitation. The motherboard already includes per-channel ADC resources used during calibration. Subject to verification of the aggregate conversion, bus, memory, and communication bandwidth, future firmware and FPGA revisions could use these resources to acquire periodic raw-waveform snapshots or potentially continuous data from all channels, with an appropriate re-design of the FPGA implementation. The system could then operate primarily in event-driven mode while temporarily enabling waveform acquisition for calibration, signal-quality assessment, artifact analysis, or selected experiments requiring complete traces.

The event-driven architecture is also a promising basis for future closed-loop systems. With a proper FPGA-level implementation, events could be accessed before packet transmission and host-side reconstruction; control decisions could therefore be executed locally using the finer internal timing rather than the 1 ms transmitted bins. Integrating programmable electrical stimulation, artifact-management circuitry, and on-device decision logic would enable the platform to detect changes in network activity and respond without continuously transferring raw data to an external computer. Low digital latency, bounded bandwidth, and compact state representations are particularly advantageous in this context.

The principal limitation of the present study remains the absence of a complete *in vitro* validation with an end-to-end platform configuration. The multichannel backend campaign bypassed the low-noise AFEs and applied a common synthetic waveform to all channels. It therefore did not reproduce channel-dependent noise, heterogeneous spike morphology, partial synchrony, electrode drift, or stimulation artifacts. Connector reliability, packet-error rate, power consumption, and long-term thermal behaviour also require systematic characterization.

Other limitations are mostly associated with the current implementation rather than with the event-driven principle. System latency was acceptable, but could be substantially reduced through the choice of a higher-speed threshold comparator component. Calibration can be improved through longer, adaptive, or quality-controlled procedures. The system also exhibited minor unintended power-supply issues. The analog and digital sections were powered by separate supplies. During preliminary bench operation, enabling the digital section was accompanied by a small decrease in the current drawn from the analog supply. This observation suggested a possible interaction between the two power domains, but the underlying current paths

were not characterized. Further investigation is required to identify the mechanism and assess whether additional power-domain isolation or controlled sequencing is needed.

Future work should therefore proceed along three connected directions. The first is an end-to-end biological validation of the complete 60-channel system over multiple cultures and long recording sessions, including comparison with reference instrumentation and assessment of threshold stability. The second is the integration of stimulation and local control algorithms to establish a closed-loop recording and actuation platform. The third is architectural translation toward larger electrode arrays and, after suitable miniaturization, power optimization, packaging, safety, and artifact management, potential *in vivo* implementations. In high-density MEA systems, the event-driven approach would not eliminate the need for scalable electronics, but it could substantially reduce the absolute communication, storage, and real-time processing burden associated with the increased channel count.

3.9 Conclusions

This chapter developed and evaluated an event-driven neural acquisition platform that converts extracellular spike threshold crossings into a compact time–channel representation. The principal system-level outcome was the substantial reduction in output data rate achieved without continuously transmitting every sampled waveform. Relative to ideal continuous 12-bit acquisition at $10\,\mathrm{kS\,s^{-1}}$, the 15-channel MCU and 60-channel FPGA implementations reduced the output payload by approximately 99.1% and 98.7%, respectively, while preserving electrode identity and millisecond-scale temporal information.

Preliminary tests of the analog front end demonstrated the acquisition of microvolt-scale artificial signals and extracellular spike-like activity with input-referred RMS noise below $10\,\mu\mathrm{V}$. The 15-channel MCU implementation established the complete acquisition workflow and achieved $F_1 \simeq 0.98$ at $\mathrm{SNR_{lin}} = 10$ for firing rates up to 300 Hz. The 60-channel FPGA implementation introduced fully parallel event monitoring and characterized the effects of SNR, threshold selection, amplitude variation, temporal clustering, and calibration conditions on detection performance.

The FPGA processing latency was approximately 80 ns, whereas the mean event-driven system latency across sessions ranged from approximately 0.34 ms to 1.45 ms under the tested conditions and was dominated by the analog threshold-detection path. Lower-SNR conditions produced the greatest latency, while the $SNR_{lin} = 20$ condition involved a different comparator-output transition and therefore requires separate physical interpretation.

The complete 60-channel prototype was realized with an estimated hardware expenditure of approximately USD 3,000, as detailed in Section 3.7.5.

The resulting architecture is particularly suited to firing-rate, burst, recruitment, synchrony, propagation, and future closed-loop applications in which event timing is more relevant than complete waveform morphology. Its main intrinsic trade-off is the loss of raw-signal information. The next development stage is therefore an end-to-end biological validation of the complete 60-channel system, followed by improvements in calibration and comparator circuitry, hybrid waveform-access modes, stimulation integration, and translation toward higher-density and potentially *in vivo* event-driven systems.

Chapter 4

Pulse-Wave Monitoring: Background

4.1 Clinical Motivation and Chapter Scope

Hypertension is a major global health burden and often progresses without evident symptoms. The World Health Organization estimated that approximately 1.4 billion adults aged 30–79 years were living with hypertension in 2024, while only slightly more than one in five had their blood pressure (BP) under control [53]. Accordingly, accessible measurement technologies could support earlier detection, treatment follow-up, and repeated assessment of cardiovascular status.

Conventional BP measurement remains essential for the diagnosis and management of hypertension. However, a cuff measurement provides a limited observation of a complex and time-varying cardiovascular system. In fact, BP does not directly characterize the mechanical properties of the arterial wall or the propagation of the pressure wave through the arterial tree. Pulse Wave analysis (PWA) and PWV provide complementary information by examining the morphology and timing of arterial pulse waves. In particular, carotid–femoral pulse wave velocity (cfPWV) has demonstrated prognostic value beyond conventional risk factors [54, 55].

The clinical relevance of these measurements does not by itself guarantee practical adoption. Established systems can provide reliable results, but they commonly require trained operators, controlled measurement conditions, and dedicated instrumentation. Extending pulse-wave assessment toward portable or wearable operation therefore requires more than sensor miniaturization. The system must acquire a mechanically weak signal from the body surface, manage the static pressure introduced by sensor

placement, preserve waveform morphology, maintain timing between anatomical sites, and provide reliable processing and quality control [56–59].

This chapter introduces the physiological and technological concepts required to interpret the two pulse-wave acquisition platforms developed in the following chapters. It first explains how arterial pressure waves are generated and how their morphology and propagation relate to vascular function. It then defines PWA, Pulse Transit Time (PTT), Pulse Arrival Time (PAT), and PWV, with particular attention to cfPWV. The chapter subsequently reviews the main sensing modalities and pressure-sensor technologies used for non-invasive pulse-wave acquisition, compares representative clinical and wearable PWV systems, and identifies the research gap addressed by the proposed platforms.

4.2 From Blood Pressure to the Arterial Pulse Wave

The systems developed in this thesis do not measure blood pressure directly. They acquire pressure-related mechanical waveforms from superficial arteries and use their morphology or relative timing to assess cardiovascular function.

4.2.1 Blood Pressure and Pulse-Wave Generation

During ventricular systole, the left ventricle ejects blood into the aorta. The resulting increase in aortic pressure distends the arterial wall and initiates a pressure disturbance that propagates through the arterial tree, as illustrated by Figure 4.1. This pulse wave represents the transmission of pressure and wall deformation rather than the displacement of the same blood volume from the heart to a peripheral site [60].

A superficial pulse waveform is not identical to intra-arterial pressure. The signal detected at the skin is affected by the tissue layers between artery and sensor, the orientation of the vessel, the contact area, and the pressure applied by the sensing interface. Nevertheless, its main temporal features can be preserved when mechanical coupling is stable and the acquisition chain has adequate bandwidth. This distinction is important for the systems presented later: the objective is not absolute pressure calibration, but reliable detection of a pressure-related waveform suitable for morphology and timing analysis.

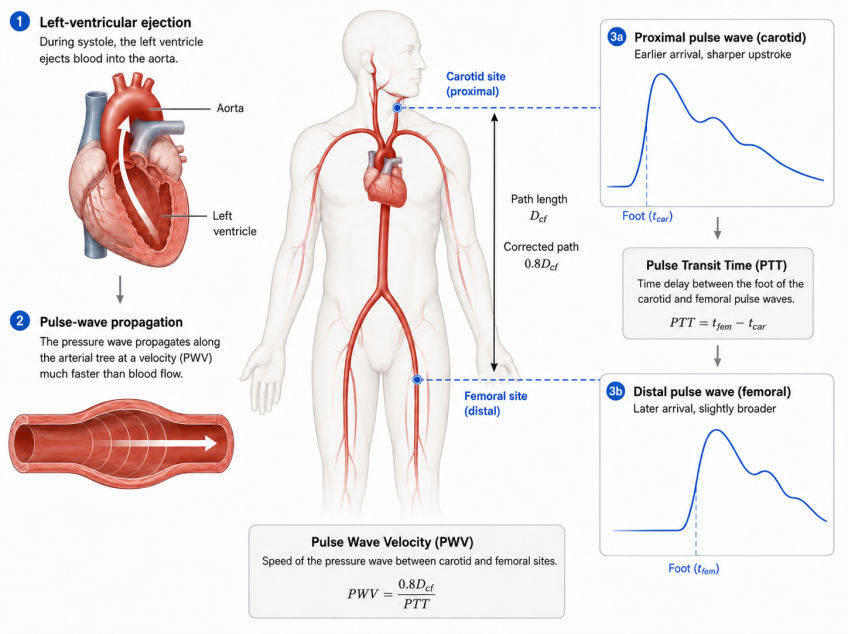


Fig. 4.1 Schematic of arterial pulse-wave generation and propagation. Left-ventricular ejection generates a pressure wave that travels along the arterial tree and is measured non-invasively at the carotid and femoral sites. Pulse transit time (PTT) is the delay between the feet of the two waveforms, and carotid–femoral pulse wave velocity (PWV) is estimated from the corrected path length and PTT.

4.2.2 Arterial Wall Mechanics and Wave Propagation

Large arteries act as compliant conduits. During systole, part of the ejected volume is accommodated by arterial expansion; during diastole, elastic recoil contributes to maintaining forward flow. This behaviour depends on the structure of the arterial wall, including elastin, collagen, smooth muscle, and calcified components. Compliance and distensibility describe the ability of a vessel to deform with pressure, whereas stiffness describes its resistance to deformation [54, 56].

Arterial stiffness generally increases with aging and vascular remodeling. A stiffer arterial segment expands less for a given pressure change and transmits the pulse wave more rapidly. Increased arterial stiffness can also alter the timing and contribution of reflected waves to the central pressure waveform and pulsatile load [54].

The relation between wall stiffness and propagation velocity can be illustrated through the simplified Moens–Korteweg expression [61, 62]:

$$PWV = \sqrt{\frac{E_{\text{inc}} h}{2\rho r}}, \tag{4.1}$$

where E_{inc} is the incremental elastic modulus of the wall, h is wall thickness, r is vessel radius, and ρ is blood density. The equation assumes an idealized thin-walled elastic tube and is not used directly in the measurements reported in this thesis. Its role is to show that, for comparable geometry and blood properties, an increase in wall stiffness produces an increase in PWV [56].

In practice, an acquired pulse wave is the result of forward propagation, local arterial mechanics, and reflected components. The early systolic region is therefore preferred for transit-time estimation because it is generally less affected by later waveform deformation and reflection than the systolic peak.

4.3 Pulse-Wave Analysis and Pulse-Wave Velocity

PWA and PWV use different information contained in the same underlying arterial phenomenon. PWA examines the shape of a pulse waveform at one or more sites, whereas PWV estimates the propagation speed of the wave between locations.

4.3.1 Information Contained in the Pulse Wave

Pulse-wave morphology reflects the combined effects of cardiac ejection and vascular load. The amplitude and slope of the systolic upstroke depend on the source waveform and local transfer conditions. The systolic peak and subsequent contour are influenced by arterial compliance, peripheral resistance, and reflected waves. The diastolic portion reflects pressure decay and runoff toward the microcirculation [63]. A schematic arterial pulse is shown in Figure 4.2.

PWA can extract amplitude-, time-, and morphology-based features. Examples include systolic and diastolic levels after calibration, pulse pressure, rise time, inflection points, reflected-wave indices, and beat-to-beat changes. The present thesis does not attempt a complete analysis of these quantities. The relevant point is that a non-invasive pressure sensor must reproduce recognizable waveform features without excessive damping, saturation, baseline drift, or distortion.

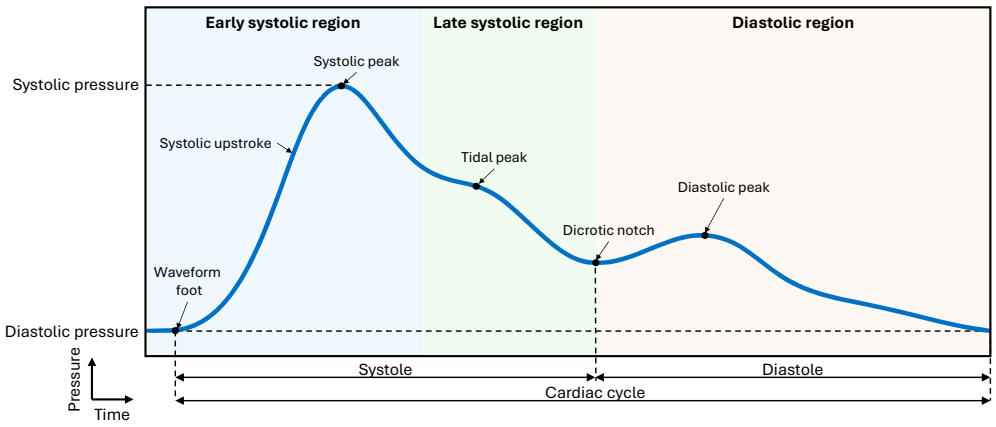


Fig. 4.2 Schematic arterial pulse waveform showing the early and late systolic regions and the diastolic phase [64]. Principal landmarks include the waveform foot, systolic upstroke and peak, tidal peak, dicrotic notch, and diastolic peak. Waveform morphology and phase boundaries are illustrative.

4.3.2 PTT, PAT, and PWV

PTT is the time required for a pulse wave to travel between two vascular measurement sites. For a proximal site A and distal site B,

$$PTT = t_B - t_A, \tag{4.2}$$

where t_A and t_B are the arrival times of corresponding pulse waves. PWV is then calculated as

$$PWV = \frac{D}{PTT}, \tag{4.3}$$

where D is the estimated propagation distance between the sites. The operational simplicity of Equation 4.3 hides two major measurement problems: defining the path length and identifying arrival times with sufficient precision [56].

PAT is related but not equivalent to PTT, and is commonly measured from an electrical cardiac reference, typically the electrocardiogram (ECG) R peak, to the arrival of a peripheral pulse. It therefore contains the pre-ejection period between ventricular electrical activation and aortic-valve opening in addition to vascular transit time. Changes in cardiac electromechanical timing can alter PAT even when vascular propagation is unchanged. PAT-based wearable systems can be useful for

tracking cardiovascular dynamics, but PAT should not be interpreted as direct PTT without accounting for this additional interval [65, 66].

Pulse-wave arrival is identified through a fiducial point rather than a directly observable physical marker. Common choices include the waveform foot, the point of maximum first derivative, the maximum of the second derivative, and the intersecting tangent point (ITP). The systolic peak is usually less suitable because its timing is more affected by waveform deformation and reflections. Foot- and tangent-based approaches are therefore commonly adopted for cfPWV measurements [54, 67, 68].

Because PWV is inversely proportional to PTT, millisecond-scale timing errors can produce relevant velocity deviations. Sensor coupling, sampling frequency, synchronization, filtering, and fiducial-point detection must therefore be considered jointly.

4.3.3 Carotid–Femoral PWV

cfPWV estimates pulse-wave propagation along the central arterial pathway and is commonly used as the reference non-invasive index of aortic stiffness [54, 67]. Pulse waves are acquired at the carotid and femoral sites. Surface measurements cannot reproduce the internal aortic path exactly. Consensus documents therefore recommend standardized distance conventions. One widely adopted expression is

$$cfPWV = \frac{0.8 D_{\mathrm{cf}}}{PTT}, \tag{4.4}$$

where D_{cf} is the direct surface distance between carotid and femoral measurement sites. The factor 0.8 provides an approximation of the travelled path that is consistent with magnetic-resonance measurements [67, 69].

Two principal acquisition strategies are used. Figure 4.3 displays a detailed comparison between the two methods. In sequential measurement, carotid and femoral waves are acquired at different times and referenced to a common signal, typically the ECG. The difference between the two ECG-to-pulse delays estimates the carotid–femoral transit time. This approach requires only one pulse-wave probe but assumes stable physiology across the sequential recordings and adds an ECG acquisition channel.

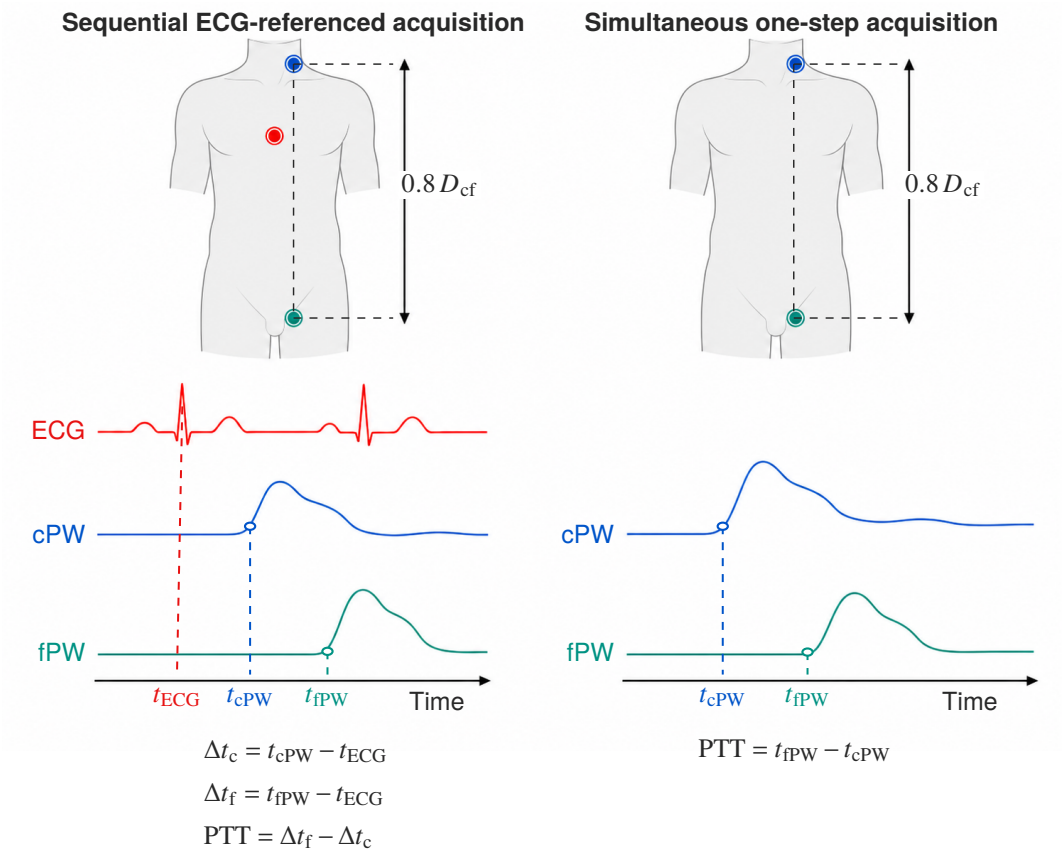


Fig. 4.3 Comparison of sequential ECG-referenced and simultaneous one-step carotid–femoral pulse wave velocity acquisition.

In simultaneous or one-step measurement, carotid and femoral waveforms are acquired during the same recording. PTT can then be measured directly from their relative delay. This avoids ECG-based realignment and reduces sensitivity to changes between sequential measurements, but it requires two sensing interfaces and a shared or accurately aligned time base.

Distance measurement, sensor placement, waveform quality, and fiducial-point choice all influence cfPWV. Accordingly, recent validation recommendations require new devices to be compared against an appropriate reference [57].

4.4 Non-Invasive Pulse-Wave Acquisition

Pulse waves can be detected through mechanical, optical, and imaging-based measurements. The selected modality determines the physical quantity observed, the accessible anatomical sites, and the extent to which the recorded signal represents local arterial pressure.

4.4.1 Sensing Modalities and Pressure-Sensor Technologies

Pressure tonometry and mechanotransducers detect the mechanical deformation generated by a superficial artery and provide a local pressure-related waveform [63, 70]. Photoplethysmography (PPG) measures pulsatile variations in optical absorption caused by changes in local blood volume [71, 72]. Doppler ultrasound and vascular imaging observe blood-flow velocity or arterial-wall motion, whereas oscillometric systems derive pulse information from pressure variations transmitted to an inflated cuff [73, 63, 74, 68]. ECG can also be combined with a peripheral pulse sensor to provide a cardiac timing reference, although the resulting interval is PAT rather than direct vascular PTT [65, 66]. Electrical bioimpedance detects pulsatile impedance changes, while electrical impedance spectroscopy (EIS) characterizes impedance as a function of frequency [75, 76]. Electrical impedance tomography (EIT) reconstructs spatial conductivity changes and has been explored for aortic pulse-timing assessment [77]. Magnetic induction tomography (MIT) maps conductivity through inductive measurements; related non-tomographic sensors have demonstrated pulse detection [78, 79]. Radar measures pulse-related displacement for non-contact assessment of pulse-wave propagation [80].

The main trade-offs are summarized in Table 4.1.

Pressure-based sensing is used to detect the local mechanical pulse at accessible carotid and femoral sites without relying on optical absorption or perfusion-dependent optical contrast. The main limitation is the strong dependence of waveform quality on anatomical localization, tissue coupling, sensor fixation, and static preload. The acquisition chain must therefore resolve a small pulsatile variation superimposed on a substantially larger placement-dependent baseline [58, 59].

The pressure transducers used for this application are commonly piezoresistive, piezoelectric, or capacitive. Their principal trade-offs are summarized in Table 4.2.

Table 4.1 Representative non-invasive sensing and timing approaches for pulse-wave and PWV assessment.

Modality	**Measured quantity**	**Main trade-off**
Pressure tonometry / mechanotransducers	Local pressure- or force-related waveform	Direct mechanical waveform acquisition, but sensitive to placement, preload, orientation, and motion [56, 68, 63, 70]
PPG	Pulsatile optical blood-volume variation	Compact and wearable-compatible, but dependent on perfusion, skin, optical coupling, and motion [71, 72]
Doppler ultrasound / vascular imaging	Blood-flow velocity or arterial-wall motion	Provides anatomical localization, but requires more complex equipment, coupling, and operator expertise [56, 73]
Oscillometric cuffs	Cuff-pressure oscillations	Simple clinical workflow, but often indirect or model-dependent and affected by cuff placement [74, 68]
ECG-referenced peripheral pulse sensing	Electrical cardiac reference combined with vascular pulse	Compact timing architecture, but measures PAT rather than direct PTT unless the pre-ejection period is accounted for [65, 66]
Electrical bioimpedance / EIS	Pulsatile impedance changes; frequency-dependent impedance in EIS	Compatible with wearable sensing, but dependent on electrode geometry, tissue contributions, and measurement frequency. Spectroscopic acquisition must preserve adequate temporal resolution [75, 76]
EIT	Spatially reconstructed conductivity changes	Enables regional extraction of vascular pulsations, but requires multiple electrodes and image reconstruction, with limited spatial resolution [77]
MIT / magnetic induction sensing	Conductivity-related inductive response; spatial reconstruction in MIT	Requires no electrode contact, but weak secondary signals and tomographic reconstruction complicate implementation. Pulse detection has also been demonstrated with non-tomographic inductive sensors [78, 79]
Radar	Pulse-related displacement detected through reflected radio-frequency signals	Enables non-contact acquisition, but requires separation of signals from different sites and rejection of unrelated body motion [80]

The platforms developed in this thesis both employ piezoresistive sensors. This choice was motivated by their simple resistance-to-voltage readout and compatibility with flexible sensing structures [59]. Their response to static and dynamic loading also supports the acquisition of pulse-related variations around a placement-dependent operating point, consistent with the baseline-adjustment strategies adopted in both platforms.

4.4.2 Instrumentation Requirements

The sensor is only one element of the measurement chain. A suitable pulse-wave system must satisfy the following coupled requirements:

- Static operating range: the sensor and readout must tolerate the placement-dependent baseline while preserving sensitivity to the pulsatile component.

Table 4.2 Qualitative comparison of pressure-sensor technologies for non-invasive pulse-wave acquisition.

Technology	Strengths and readout	Main concerns
Piezoresistive	Simple resistance-to-voltage conversion; compact and compatible with flexible substrates	Preload dependence, drift, hysteresis, nonlinearity, and temperature sensitivity [58, 59]
Piezoelectric	High dynamic sensitivity; charge or high-impedance voltage readout	Leakage, shielding, and package-motion sensitivity [58, 59, 81]
Capacitive	High low-pressure sensitivity and low static power	Parasitic capacitance, humidity, environmental coupling, and more complex readout [59, 81]

- Waveform preservation: the acquisition chain must preserve the early systolic upstroke and the features used for morphology or fiducial-point detection.
- Temporal resolution and synchronization: sampling and inter-channel alignment must support millisecond-scale transit-time estimation.
- Wearable integration: for wearable devices size, power, wireless communication, comfort, and data handling must be considered jointly.

High sensor sensitivity alone is therefore insufficient. Performance depends on the interaction between mechanical coupling, sensor operating point, analog conditioning, timing, packaging, and signal processing. Conversely, a sensor with moderate intrinsic sensitivity can provide reliable pulse-wave acquisition when integrated within a properly designed measurement chain.

4.5 State of the Art in PWV Measurement Systems

PWV systems differ primarily in how pulse waves are acquired, how the transit time is determined, and which anatomical distance is used. These choices affect operator dependence, measurement directness, portability, and comparability between devices. Established clinical systems generally prioritize standardized acquisition and signal-quality control, whereas portable and wearable approaches attempt to reduce equipment size and enable repeated measurements.

4.5.1 Clinical Reference and Commercial Systems

Tonometric systems such as SphygmoCor and PulsePen acquire local arterial pressure-related waveforms and commonly use ECG-referenced sequential measurements, although cuff-assisted or dual-probe configurations are also available. Carotid–femoral PWV measured by applanation tonometry is widely regarded as the gold-standard non-invasive assessment of arterial stiffness, and SphygmoCor is among the most established reference platforms implementing this approach [54, 68]. Tonometric measurements provide direct access to pulse-wave morphology but remain sensitive to arterial localization, contact pressure, and operator technique [82, 68].

Complior instead uses two external mechanotransducers to acquire proximal and distal pulse waves simultaneously. This removes physiological variation between sequential measurements and provides direct dual-site PTT estimation, but requires stable positioning and coupling of both sensors throughout the acquisition [83].

Cuff-based systems reduce the need for precise point localization. Vicorder acquires pulse-related signals from two cuff locations and estimates transit time from simultaneous recordings. Arteriograph derives aortic stiffness parameters from an upper-arm oscillometric waveform and an internal model. These approaches simplify the measurement procedure, but their results depend more strongly on cuff positioning, path-length convention, and device-specific processing assumptions [84, 74].

Table 4.3 summarizes these representative implementations.

Table 4.3 Representative systems for PWV and arterial-stiffness assessment.

System	Sensing principle	Acquisition strategy	Main trade-off
SphygmoCor	Tonometry and cuff-based sensing, depending on configuration	Sequential ECG-referenced or hybrid cfPWV measurement	Established reference use; trained placement and controlled acquisition required [54, 68]
Complior	Dual mechanotransducers	Simultaneous two-site measurement	Direct PTT; stable coupling required at both sites [83]
PulsePen	Tonometry with ECG or dual probes	Sequential or two-probe measurement	Portable and waveform-based; operator-dependent contact [82]
Vicorder	Cuff-based mechanosensing	Simultaneous two-site measurement	Simplified setup; cuff and path-length dependence [84]
Arteriograph	Upper-arm oscillometry	Model-based aortic PWV estimation	Simple workflow; indirect rather than dual-site cfPWV measurement [74]

4.5.2 Portable and Wearable Approaches

Wearable cardiovascular systems frequently combine ECG and PPG because both signals can be acquired using compact, low-power hardware. Such systems measure PAT, which includes the pre-ejection period, rather than transit time between two vascular sites. They are therefore widely investigated for cuffless blood-pressure monitoring but do not directly measure vascular PTT or cfPWV [65, 66, 85].

Dual-site PPG estimates PTT from the delay between two optical pulse waves, thereby excluding the pre-ejection period. Most implementations, however, assess local or peripheral arterial paths and remain sensitive to perfusion, optical coupling, skin properties, and motion [86, 72].

Pressure-based wearable systems are more closely aligned with the conventional foot-to-foot method because they acquire local mechanical pulse waves directly from superficial arteries. Their implementation ranges from individual flexible sensors to multi-element arrays designed to reduce localization sensitivity [87, 58, 88–90].

Alternative approaches employ Doppler ultrasound, bioimpedance, seismocardiography, or phonocardiography to detect pulse propagation or provide additional timing references [73, 91–93].

The central challenge is therefore not simply detecting a periodic vascular signal, but acquiring two anatomically meaningful waveforms with stable coupling and preserved relative timing.

4.5.3 Current Technological Limitations

Three main limitations recur across portable and wearable PWV systems. First, many approaches measure PAT or peripheral PTT rather than cfPWV, limiting direct comparison with central arterial-stiffness measurements. Second, direct carotid–femoral acquisition remains dependent on stable dual-site coupling: sequential methods are exposed to inter-recording physiological variation, whereas simultaneous methods require accurate inter-channel synchronization. Third, wireless implementation must preserve timing despite independent clocks, packet-based communication, and power constraints.

These limitations are coupled rather than independent. Increased fixation may improve contact while shifting the sensor operating point; reduced sensor size may improve wearability while increasing localization sensitivity; and wireless operation may eliminate cables while introducing synchronization and communication constraints. The relevant objective is therefore the integration and validation of the complete measurement chain rather than the optimization of an isolated sensor parameter.

4.6 Research Gap and Objectives

Previous systems, including Athos and the wireless platform developed by Valerio et al., demonstrated simultaneous carotid–femoral pulse-wave acquisition and its integration with dedicated electronics and software [87, 89]. These developments provided the starting point for the present research. Stable acquisition nevertheless remains sensitive to sensor positioning, contact pressure, and mechanical coupling [56, 58]. The practical problem motivating this work was how to address these factors jointly within a simple acquisition setup, while reducing the need for continuous manual probe handling.

The objective was therefore to investigate the complete pressure-wave acquisition chain, from the sensing element and mechanical interface to electronic conditioning, synchronized recording, and host software. The guiding research question is: How can these elements be integrated to improve pulse-wave acquisition quality and reduce dependence on manual probe handling within a simple setup for carotid–femoral measurements?

The initial investigation explored a custom graphene-based piezoresistive sensor, motivated by the pressure sensitivity and tunable mechanical response reported for functionalized graphene materials [94]. Chapter 5 describes the development of the sensing structure, mechanical interface, calibrated readout, and acquisition software, together with controlled tests and preliminary physiological recordings [5].

Despite promising outcomes, the mechanical robustness and electrical constraints encountered with this custom sensor subsequently motivated a second development stage. Commercial force-sensing resistors were selected for their low cost, flexibility, small form factor, and ease of integration. Chapter 6 presents the resulting wearable

platform, combining baseline-adaptive conditioning, synchronized acquisition, and wireless communication. Wearable fixation was introduced to reduce the need for continuous manual stabilization, while retaining operator involvement in positioning and signal-quality verification. The resulting cfPWV estimates were compared with SphygmoCor under supervised clinical conditions [4].

Chapter 5

Graphene-Based Pressure Sensing Platform for Pulse Wave Acquisition

5.1 Introduction and Motivation

The previous chapter introduced the physiological and technological background of cardiovascular pulse-wave monitoring, with particular attention to the role of pressure-related arterial waveforms in the assessment of vascular function. Building on that background, this chapter presents the development of a custom reduced-graphene-oxide pressure sensing platform for non-invasive pulse-wave acquisition, based primarily on the work reported in [5].

The acquisition of arterial pulse waves from superficial anatomical sites is a challenging task from an instrumentation perspective. The useful signal is generated by small local pressure variations associated with arterial pulsation, but it is strongly affected by sensor placement, static contact pressure, mechanical coupling with the skin, and the properties of the sensing element [68, 58, 95]. A practical pressure sensing platform for this application must therefore combine a sensitive transducer, a mechanically compliant structure, a stable pressure-transfer interface, and precise electronic readout circuit.

As discussed in Section 4.4.1 of the previous chapter, piezoresistive pressure sensors represent a suitable solution for this type of application because they convert mechanical deformation into resistance variations using relatively simple electronic

conditioning circuits. Among the possible active materials, graphene-derived materials are attractive because their electrical properties can be modulated by mechanical deformation and because they can be processed into flexible sensing structures. In this work, the active material is a modified reduced-graphene oxide functionalized with triethylene glycol amine (rGO-R1). The material acts as the pressure-sensitive layer of the sensor and provides resistance modulation when the flexible structure is deformed.

Previous work demonstrated the low-pressure sensitivity, tunable mechanical response, and heartbeat-detection capability of functionalized rGO [94]. Material sensitivity alone, however, does not establish the dynamic behaviour of an acquisition platform, which also depends on sensor geometry, contact preload, mechanical coupling, and electronic readout [58, 95]. The specific issue addressed in this chapter is therefore how to translate these material properties into a flexible-PCB sensing structure and a complete acquisition chain suitable for recording small arterial pulse variations under contact preload.

The objective of this work is therefore to assess whether an rGO-R1-based flexible pressure sensor, combined with dedicated mechanical and electronic interfacing, can operate as a proof-of-concept platform for non-invasive arterial pulse-wave acquisition. The focus is therefore on the development and preliminary characterization of the sensing platform, including its pressure-dependent response, dynamic behaviour, and ability to acquire physiological pulse-related waveforms from superficial arterial sites. These measurements are intended to demonstrate acquisition capability and waveform quality, while more structured validation and pulse wave velocity (PWV) estimation are addressed by the wearable system presented in the following chapter.

Sections 5.2–5.6 describe the platform design, from the sensor and mechanical holder to the electronics, firmware, and host software. Sections 5.7 and 5.8 present the validation methodology and results, followed by the discussion and conclusions.

5.2 Design Strategy and System Overview

The rGO-based pressure sensing platform was designed to acquire pressure-related arterial pulse waveforms from superficial anatomical sites. In this application, the sensing system must detect small dynamic pressure variations while operating under

a static contact pressure introduced by probe positioning. The design strategy therefore required the joint optimization of the sensing element, mechanical interface, electronic readout, and acquisition workflow.

5.2.1 Pulse Wave Acquisition Requirements

The platform was intended to address three main acquisition requirements:

- Sensitivity and waveform preservation. The sensing and conditioning chain must detect small pulse-related pressure variations. As an indicative example, a flexible piezoresistive sensor positioned over the radial artery recorded a pressure variation of approximately 170 Pa [96]. The pressure transmitted to the sensor depends on anatomical site and mechanical coupling, so this value provides only an indicative scale. The sensor must therefore resolve small pulsatile variations while accommodating the static contact preload. The acquisition bandwidth must preserve the cardiac fundamental and relevant harmonics, with the main pulse-wave components typically below 20 Hz [97].
- Stable mechanical coupling and usability. The probe must support repeatable positioning, allow adjustment of contact pressure, and remain practical and comfortable during acquisition. The readout must accommodate the placement-dependent baseline while retaining sensitivity to the superimposed pulsatile component.
- Acquisition and data handling. The system must support simultaneous two-site recording with millisecond-scale sampling resolution for subsequent pulse-timing analysis. A sampling interval no greater than 1 ms, corresponding to at least 1 kHz per channel, provides this temporal grid directly and is consistent with reference acquisition protocols [98]. Real-time visualization and storage of both waveforms are also required.

5.2.2 System Overview

Figure 5.1 summarizes the platform architecture. The system comprises up to two identical pen-shaped probes, an external embedded module referred to as the main board, and a host computer with a graphical user interface (GUI). Each probe

Table 5.1 Main specifications of the rGO-R1 pulse-wave acquisition platform [5].

Specification	Value
Sensing element	rGO-R1 cross-shaped flexible-PCB sensor
Sensor baseline resistance	Approximately 1 Ω to 10 Ω
Sensor interface	Current-based Wheatstone bridge
Automatic balancing	Dual DAC-controlled offset voltages
Current-to-voltage gain	$GR_F = 690\,\mathrm{V\,A^{-1}}$
Analog passband	0.1 Hz to 20 Hz
Embedded platform	Pyboard SF2W D-series
ADC resolution	12-bit
Sampling frequency	1 kHz per channel
Communication	USB serial link
Host software	Python/Kivy GUI
Estimated current consumption	> 100 mA per probe

integrates the rGO-R1 sensor and its local readout and calibration circuitry, whereas the main board performs further conditioning, digitization, and communication. The host software controls acquisition, displays both channels, and stores the recorded data. The principal system specifications are summarized in Table 5.1.

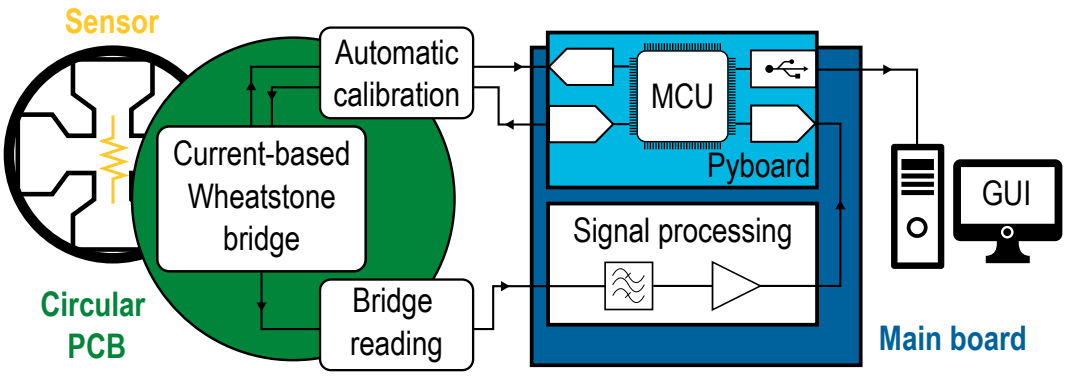


Fig. 5.1 The pressure sensor is interfaced to the circular PCB, which performs signal readout; the external embedded module, referred to as the main board, performs further signal processing, digitization, and communication; and the host software running on a computer provides real-time visualization and data storage through its GUI.

5.3 Sensor Design

The pressure sensor combines the rGO-R1 active material, a flexible PCB substrate, and a geometry designed to enhance deformation in the sensing regions. Material preparation and deposition were developed in collaboration with the Nanochemistry Laboratory at the University of Strasbourg, whereas the sensor geometry and system integration were engineered for pulse-wave acquisition.

5.3.1 Material and Working Principle

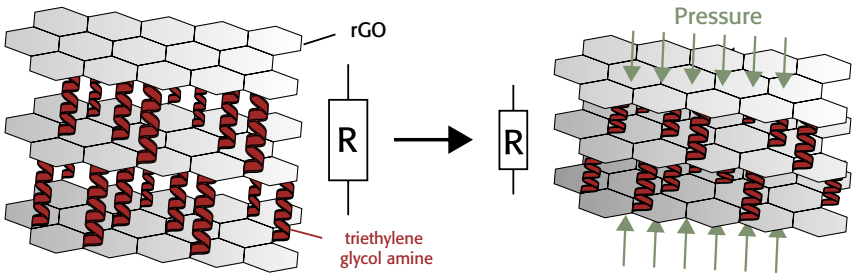
The active material is rGO-R1, a reduced graphene oxide functionalized with triethylene glycol amine. In the sensor, the material is deposited on opposing flexible-PCB surfaces that form the pressure-sensitive electrical interface. Its piezoresistive response is attributed to pressure-induced changes in the contact area, separation, and conductive pathways between graphene-based regions.

When pressure is applied, deformation and compression of the assembled structure increase the effective electrical contact within the conductive network, producing a decrease in sensor resistance. Unloading allows partial recovery of the initial geometry and resistance. The response therefore depends not only on the material, but also on coating coverage, substrate geometry, preload, and mechanical recovery. The device is consequently used as a pressure-related waveform sensor rather than as an intrinsically calibrated absolute-pressure transducer [5].

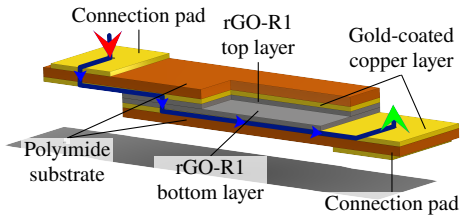
The detailed chemical synthesis and material-level optimization are reported elsewhere [94]; here the focus is the integration of rGO-R1 into a custom sensor geometry, mechanical interface, and acquisition chain.

5.3.2 Sensor Structure and Geometry Optimization

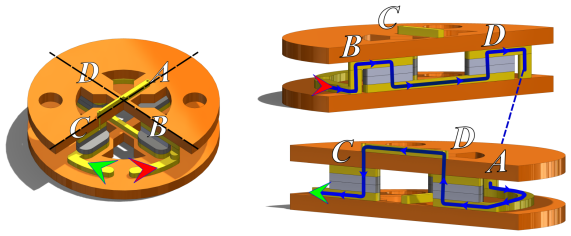
A polyimide flexible PCB was adopted as both mechanical support and electrical interface. Polyimide provides flexibility, durability, and compatibility with standard manufacturing, enabling reproducible geometries, conductive paths, and contact pads [99]. Gold-finished regions were used for rGO-R1 deposition, while pads on the opposite side interfaced with the readout electronics. The active material was



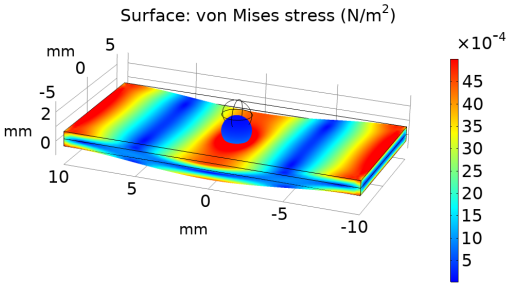
(a) rGO-R1 piezoresistive mechanism



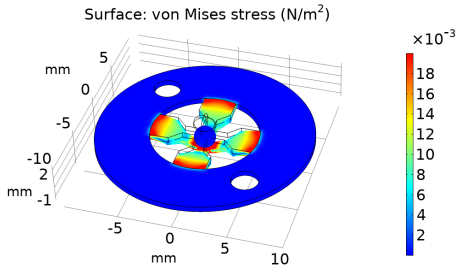
(b) Rectangular sensor section



(c) Cross-shaped sensor section



(d) Rectangular-sensor stress distribution



(e) Cross-shaped-sensor stress distribution

Fig. 5.2 Working principle and geometry of the custom rGO-R1 flexible-PCB sensors. (a) Conceptual piezoresistive mechanism, in which applied pressure increases conductive contact and reduces resistance. (b) Rectangular sensor section with labeled polyimide, metal, and rGO-R1 layers; current path is shown with a blue line; current inlet and outlet are indicated with red and green arrows. (c) Cross-shaped sensor section using the same color code; terminals A–D are indicated and the illustrated current path extends from B to C. (d,e) von Mises stress distributions for the rectangular and cross-shaped geometries under a central 0.5 N load applied through a 2 mm rounded tip.

spray coated through a shadow mask that defined the sensing regions [5, 94, 100], after which two coated flexible PCBs were assembled face-to-face.

Two sensor geometries were considered: a rectangular shape, used as a reference configuration, and a cross-shaped geometry, introduced to improve stress concentration while maintaining a compact design. The two layouts differ in the electrical path

defined by the coated regions, as illustrated by Figure 5.2b and Figure 5.2c. In the cross-shaped sensor, the intended current path traverses four coated regions before reaching the output terminal. The final cross-shaped devices used in the probes exhibited baseline resistances of approximately 1 Ω to 10 Ω.

Finite-element simulations compared the two geometries under the same central 0.5 N load applied through a 2 mm rounded tip. The rectangular sensor was constrained at its extremities, whereas the circular crown of the cross-shaped sensor was fixed. As shown in Figures 5.2d and 5.2e, the cross-shaped geometry concentrated stress near the center and arm-to-crown transitions, with peak values more than one order of magnitude higher than in the rectangular configuration. This result motivated its selection for the final probe, since concentrating deformation in the coated regions was expected to improve conversion of small pressure variations into measurable resistance changes.

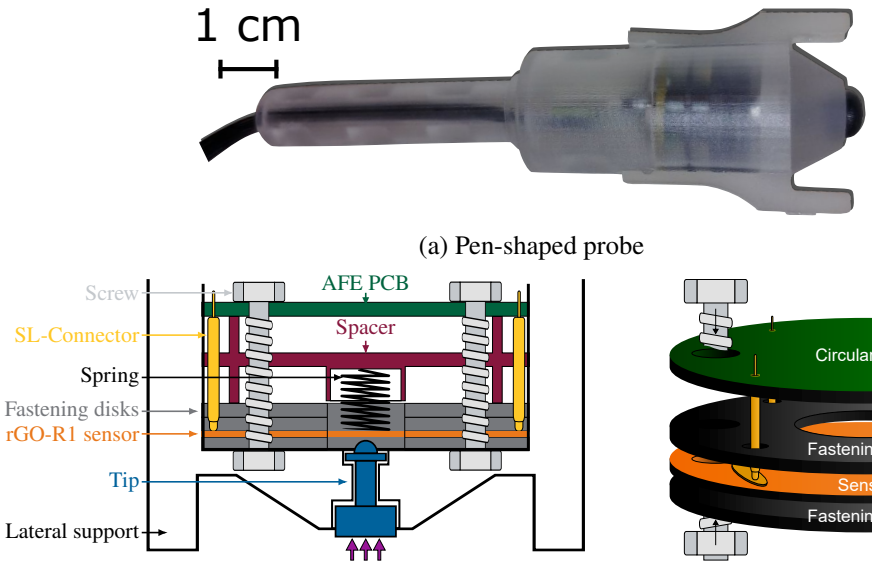
5.4 Mechanical Integration and Sensor Coupling

The mechanical integration of the pressure sensor plays a central role in arterial pulse-wave acquisition. Since the sensor response depends on the deformation of the rGO-R1 active layer, the interface must transfer pressure variations from the anatomical site to the sensing structure in a controlled and repeatable way. For this reason, a dedicated pen-shaped holder was developed around the cross-shaped sensor geometry, with the aim of supporting manual positioning, pressure transfer, electrical connection, and mechanical recovery during acquisition. The two resulting probes used in the dual-site configuration are mechanically identical. The complete mechanical assembly is shown in Figure 5.3.

5.4.1 Pen-Shaped Holder and Pressure-Transfer Interface

The holder consists of an ergonomic body and a removable distal head. The body provides a stable grip, houses the sensor–analog-front-end assembly in an internal recess, and routes the connection wires to the main board.

The head contains an axially guided pressure-transfer tip composed of a chamfered cylindrical skin-contact element and a proximal shaft ending in a hemispherical



(a) Pen-shaped probe

(b) Cross-sectional assembly

(c) Sensor–PCB interface

Fig. 5.3 Mechanical integration of the cross-shaped sensor. (a) Photograph of the assembled pen-shaped probe. (b) Cross-sectional view showing the skin-contact tip, pressure-transfer shaft, lateral supports, fastening disks, elastic-return spring, spacer, circular analog-front-end PCB, and spring-loaded electrical connectors. (c) Three-dimensional rendering of the detachable sensor-to-PCB interface.

surface that loads the center of the sensing diaphragm. Lateral supports distribute part of the applied preload over a wider area, reducing excessive diaphragm compression and improving comfort. Although the removable head permits alternative pressure-transfer geometries, only the configuration described here was evaluated.

5.4.2 Sensor Mounting and Electrical Coupling

As described in Section 5.3.2, the cross-shaped sensor is obtained by assembling two rGO-R1-coated flexible PCBs in a face-to-face configuration. This assembly was implemented using two 3D-printed fastening disks positioned on opposite sides of the sensor. The disks keep the two flexible PCBs aligned and fixed, leaving the central sensing region accessible to the pressure-transfer shaft.

The fastening disks constrain the external crown to the circular analog-front-end PCB through four screws while leaving the central region free to deform, as shown in Figure 5.3c. Spring-loaded connectors provide detachable electrical contact with the sensor pads and accommodate assembly tolerances. A 3D-printed spacer

compensates for connector compression and maintains a fixed sensor–PCB distance, preserving mechanical alignment and reliable contact.

5.4.3 Spring-Assisted Elastic Recovery

During preliminary pulse-wave acquisitions, the flexible sensor structure showed slow and incomplete elastic recovery after the application of relatively high static pressure. This behaviour reduced the ability of the sensing diaphragm to return to its initial configuration and decreased the sensitivity to small pressure variations.

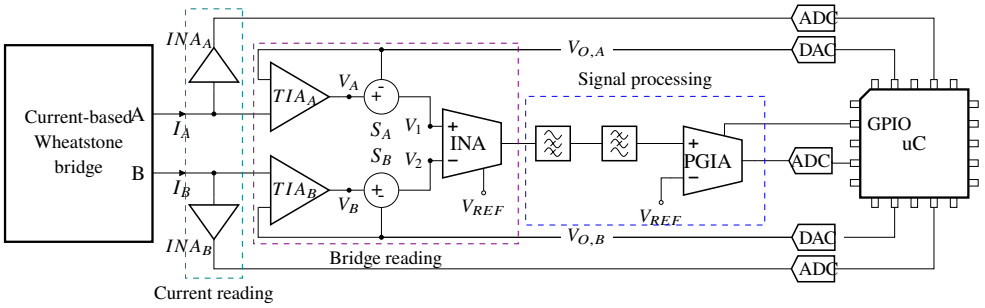
To mitigate this effect, a spring with a nominal stiffness of $0.17\,\mathrm{N\,mm^{-1}}$ was integrated into the mechanical assembly and centred by a guide in the spacer. Its free length was chosen to match the available distance to the undeformed sensor, without applying nominal preload at rest. The spring was not attached to the sensor and exerted a restoring force only when compressed.

The stiffness was selected empirically through preliminary actuator-driven tests using simplified pulse-like mechanical inputs. The acquired waveforms were visually compared without a spring and with springs of different stiffness. The selected spring appeared to improve elastic recovery without obvious additional waveform distortion under the tested conditions. These comparisons were qualitative and did not quantify the spring contribution to the dynamic response.

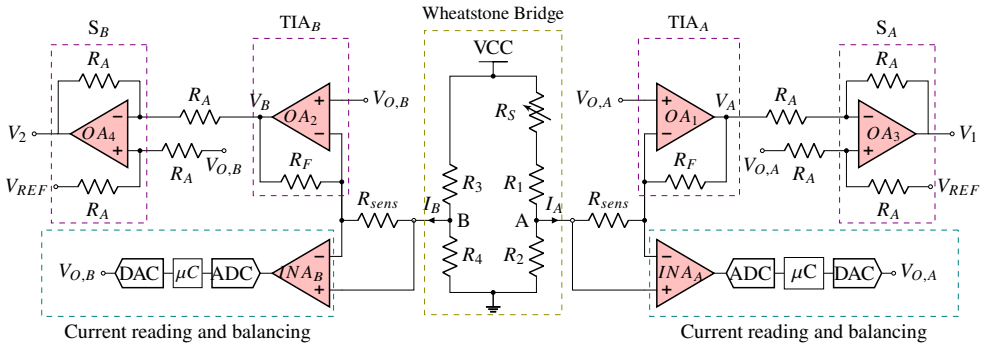
5.5 Electronic Acquisition System

The electronic acquisition system was designed to convert the pressure-induced resistance variations of the rGO-R1 sensors into voltage signals suitable for digital acquisition and visualization. Since the expected resistance variations are small, the acquisition chain provides sensor calibration, resistance-to-voltage conversion, analog conditioning, and continuous dual-channel acquisition [5].

Figure 5.4a presents a system-level overview of the acquisition chain, while Figure 5.4b shows the bridge readout and balancing circuit.



(a) Acquisition chain



(b) Current-based Wheatstone bridge, readout and balancing circuit

Fig. 5.4 Electronic acquisition architecture. (a) Simplified block diagram of the overall acquisition chain. The probe-side circuitry converts the sensor resistance variation into a differential voltage, the main board provides filtering and programmable gain, and the MCU controls bridge balancing, digitizes the two channels, and communicates with the host computer. (b) Detailed schematic of the current-based Wheatstone bridge, readout, and balancing circuit.

5.5.1 Current-Based Wheatstone Bridge

The first readout stage was based on a current-based Wheatstone bridge adapted from the architecture proposed in [101]. A conventional Wheatstone bridge detects resistance variations through the differential voltage between its two central nodes. For very small resistance changes, however, the useful voltage can be limited by the initial bridge offset caused by component mismatch and by the nominal sensor resistance, thereby restricting the gain that can be applied in subsequent stages. The current-based architecture instead measures the differential currents at the bridge nodes and converts them through transimpedance stages. This approach retains differential common-mode rejection, increases the achievable detection sensitivity, and allows the initial bridge imbalance to be compensated electronically without

modifying the bridge elements. The reference architecture reported a resistance-variation resolution of approximately 0.001 Ω, with substantially higher detection sensitivity than a basic voltage-output bridge.

As shown in Figure 5.4b, the bridge was supplied between $V_{CC} = 3.3\,V$ and ground. The sensor resistance R_S is placed in series with one bridge resistor and acts as the only variable resistive element. Pressure-induced deformation of the rGO-R1 layer changes R_S and modifies the currents flowing through the two bridge branches. The useful electrical quantity is therefore:

$$I_{OUT} = I_B - I_A. \tag{5.1}$$

Because R_S has a non-zero baseline value, the uncompensated bridge produces a non-zero current difference even in the absence of a pressure variation. Assuming equal bridge resistors $R_1 = R_2 = R_3 = R_4 = R_0$, this current difference can be expressed as:

$$I_{OUT} = \frac{V_{CC}}{2} \frac{R_S}{R_0(R_0 + R_S)}. \tag{5.2}$$

The baseline contribution is therefore compensated through an electronic calibration procedure before acquisition, as described in Section 5.5.3. Once this operating point has been established, subsequent pressure-induced variations of R_S generate the differential current processed by the acquisition chain.

In this implementation, R_0 was set to 30 Ω. This value was selected as a compromise between sensitivity, operating range, and electrical loading. The cross-shaped sensor has a baseline resistance close to 1 Ω, with possible values extending up to approximately 10 Ω. A much larger R_0 would reduce sensitivity to small variations of R_S, whereas a much smaller value would increase branch current, power consumption, and resistor heating. With $R_0 = 30\,\Omega$, $V_{CC} = 3.3\,V$, and R_S between approximately 1 Ω to 10 Ω, the nominal bridge current was approximately 102 mA to 109 mA per probe. The bridge therefore dominated the overall power consumption of the platform.

5.5.2 Analog Conditioning Chain

The analog conditioning chain converts the current imbalance generated by the bridge into a voltage waveform suitable for digital acquisition. The first stage consists of

transimpedance amplifiers (TIAs), denoted as TIA_A and TIA_B, that convert the branch currents I_A and I_B into voltage signals. Their output depends on the feedback resistance R_F and on the calibration voltages applied during bridge balancing.

After current-to-voltage conversion, the calibration contribution is removed from both branches by summing amplifiers (S_A, S_B). The resulting signals (V_1 and V_2) are referred to mid-supply and then subtracted and amplified, producing:

$$V_{\mathrm{OUT}} = G\,(V_2 - V_1) + \frac{V_{\mathrm{CC}}}{2} = G R_F I_{\mathrm{OUT}} + \frac{V_{\mathrm{CC}}}{2}, \tag{5.3}$$

where R_F is the transimpedance feedback resistance and G is the gain of the differential amplification stage. In the implemented circuit, $R_F = 7.5\,\Omega$, $G = 92$, and $V_{\mathrm{CC}} = 3.3\,\mathrm{V}$, corresponding to a nominal current-to-voltage gain GR_F of $690\,\mathrm{V\,A^{-1}}$ and an output reference of $1.65\,\mathrm{V}$.

The voltage generated on the circular PCB is routed to the main board for further processing. This stage includes third-order Sallen–Key high-pass and low-pass filters, with cutoff frequencies of 0.1 Hz and 20 Hz, respectively. The 0.1 Hz high-pass cutoff was selected to attenuate slow variations in contact pressure while retaining beat-scale waveform components, whereas the 20 Hz low-pass cutoff limits high-frequency noise while preserving the principal morphology of arterial pulse waves [63, 58].

The final analog stage is a programmable-gain INA, allowing the signal amplitude to be adapted to different sensors, contact conditions, and anatomical sites. The conditioned outputs are then acquired by the ADCs on the MCU.

5.5.3 Automatic Bridge Balancing

The non-zero sensor baseline, bridge mismatch, and slow drift produce a residual bridge current before acquisition. If this contribution is not compensated, the following amplification stages may operate outside their optimal range or saturate before the dynamic pulse component is acquired. The readout circuit therefore performs an automatic calibration before each acquisition.

The adopted circuit uses a voltage-mode balancing strategy. The two TIAs are connected to the bridge output nodes through their inverting inputs, while programmable calibration voltages $V_{O,A}$ and $V_{O,B}$ are applied to their non-inverting inputs. Under negative feedback, the bridge nodes are forced toward these voltages.

Adjusting $V_{O,A}$ and $V_{O,B}$ therefore changes the operating point of the two branches and reduces the residual currents I_A and I_B toward the selected null condition.

The calibration voltages are generated by DACs on the circular PCB and controlled by the MCU on the main board. During calibration, the residual current is estimated from the voltage drop across R_{sense} after amplification by an INA. The MCU uses this measurement to update $V_{O,A}$ and $V_{O,B}$ according to the routine described in Section 5.6.2. Once calibration is complete, the selected values are maintained during acquisition and their contribution is removed by the following summing stages.

5.5.4 Embedded Acquisition Module

The main board integrates the remaining analog-processing stages and a pyboard SF2W D-series based on an STM32 microcontroller [102]. The MCU controls bridge calibration, programmable-gain selection, analog-to-digital conversion, and communication with the host computer. Each probe channel is connected to a dedicated 12-bit ADC integrated in the STM32 MCU. Acquisition is controlled by an internal timer operating at 1 kHz. At each timer interrupt, one sample is acquired from each connected channel and the resulting sample pair is transmitted to the host computer through the USB serial link. The two channels are therefore synchronously sampled at each 1 ms acquisition interval, supporting simultaneous dual-site pulse-wave recording.

5.6 Firmware and Host Software

The platform is coordinated by embedded firmware running on the STM32-based pyboard and by host software running on the acquisition computer. The firmware manages hardware initialization, bridge calibration, programmable-gain selection, synchronous dual-channel sampling, and USB communication. The host software implements a GUI for acquisition control, real-time visualization, and data storage.

5.6.1 Firmware Workflow

The firmware was implemented in MicroPython and manages hardware initialization, bridge calibration, programmable-gain selection, and timer-driven acquisition. During recording, the 1 kHz timer callback acquires one 12-bit sample from each probe channel and transmits the paired samples to the host computer through USB.

5.6.2 Bridge-Balancing Routine

The bridge-balancing routine implements the firmware control of the analog calibration circuit described in Section 5.5.3. Its purpose is to reduce the residual bridge current before acquisition so that the amplification stages can operate around a suitable baseline.

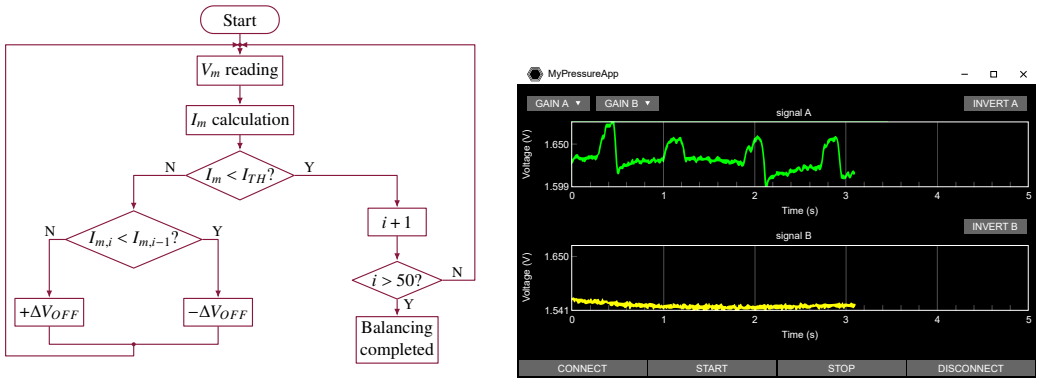
At the beginning of the routine, the two calibration voltages $V_{O,A}$ and $V_{O,B}$ are initialized at $V_{CC}/2$. The MCU then programs the corresponding DAC and estimates the residual current from the amplified voltage generated across R_{sense}.

The measured current I_m is compared with a predefined threshold I_{TH}. If I_m exceeds I_{TH}, the corresponding calibration voltage is updated by a step ΔV_{OFF} in the direction that reduces the current, and the measurement is repeated. The same procedure is then applied to the second branch. The balancing threshold was set to $I_{\mathrm{TH}} = 1\,\mathrm{mA}$, and the calibration-voltage update step was set to $\Delta V_{\mathrm{OFF}} = 15\,\mathrm{mV}$. A branch was considered balanced when I_m remained below I_{TH} for 50 consecutive iterations. These parameters were selected empirically as a compromise between the residual current and calibration duration and stability.

After both branches have been balanced, the selected calibration voltages are maintained during acquisition. A flowchart of the routine is illustrated in Figure 5.5a.

5.6.3 Host Software and User Interface

The host software was developed in Python using the Kivy framework and implements the GUI shown in Figure 5.5b. It controls connection, calibration, gain selection, and acquisition, while displaying the two channels and storing the received sample pairs for offline processing.



(a) Bridge-balancing routine (b) Host-software user interface

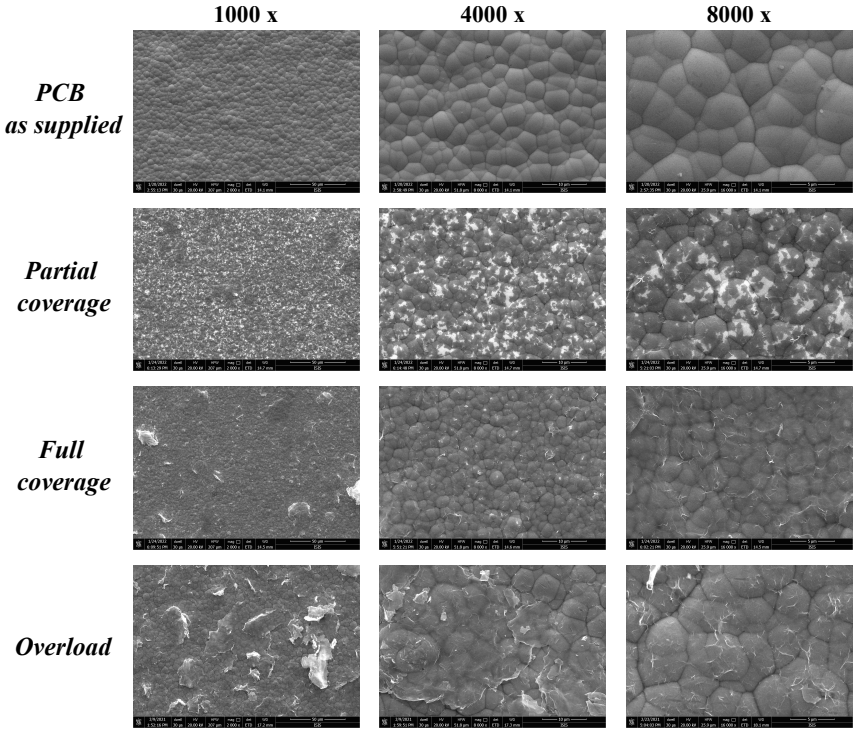
Fig. 5.5 Bridge balancing algorithm (a) and host-software UI (b) for the rGO-R1 platform. The balancing routine iteratively updates $V_{O,A}$ and $V_{O,B}$ from the measured residual currents, while the UI provides connection, calibration, gain selection, two-channel visualization, acquisition control, and data storage.

5.7 Experimental Validation Methodology

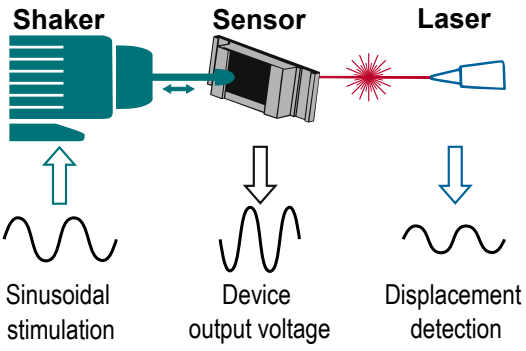
The validation was organized at three levels of integration. First, sensor-level static and cyclic tests characterized the pressure-dependent resistance response of the rGO-R1-coated flexible PCB and its stability under repeated loading. Dynamic mechanical tests then assessed waveform preservation and signal quality under controlled time-varying stimulation. Finally, preliminary physiological acquisitions evaluated the operation of the complete platform on superficial arterial sites. The corresponding results are reported in Section 5.8.

5.7.1 Sensor-Level Characterization

Sensor-level characterization was performed in collaboration with the Nanochemistry Laboratory at the University of Strasbourg to evaluate the piezoresistive response of rGO-R1-coated flexible PCB sensors and the stability of this response under repeated loading. The tests were performed on rectangular sensors matching the geometry, electrode layout, and rGO-R1 coverage of the previous indium-tin-oxide-coated polyethylene terephthalate (ITO–PET) implementation [94], thereby allowing the effect of the substrate change to be examined. The previous ITO–PET device reported a sensitivity of $0.82\,\mathrm{kPa}^{-1}$.



(a) Coating morphology and coverage classes



(b) Dynamic mechanical test bench



(c) Physiological acquisition setup

Fig. 5.6 Experimental validation at sensor, system, and physiological levels. (a) SEM images used to classify the rGO-R1 coating. (b) Shaker–vibrometer setup used to compare the electrical output with controlled diaphragm displacement. (c) Carotid pulse-wave acquisition setup.

For these characterizations, sensors with different levels of coverage of the deposited rGO-R1 layer were used. The morphology and surface coverage were inspected by scanning electron microscopy (SEM), as shown in Figure 5.6a. Based on the observed coverage, the samples were classified as having partial coverage, full

coverage, or overload. This classification was used to compare how the amount and distribution of active material affected both sensitivity and durability.

Static piezoresistive characterization was performed using a digital force gauge mounted on a motorized test stand. A circular compression plate with a diameter of 1.15 cm was used to apply controlled loads to the sensor surface. The applied force was set from 0.005 N to 0.2 N, corresponding to an estimated pressure range of approximately 0.05 kPa to 2 kPa, selected to characterize the low-pressure response of the material. Because the effective contact area and static preload of the final probe were not measured *in vivo*, it should not be interpreted as an exact representation of the pressure applied during acquisition. During each loading condition, the electrical resistance was measured using a source-measure unit.

The sensor response was expressed as relative resistance variation:

$$\frac{\Delta R}{R_i} = \left| \frac{R(P) - R_i}{R_i} \right|, \tag{5.4}$$

where R_i is the initial resistance before loading and $R(P)$ is the resistance measured at the applied pressure P. The pressure sensitivity was then estimated from the slope of the relative resistance variation as a function of pressure in the low-pressure range:

$$S = \left| \frac{\partial(\Delta R/R_i)}{\partial P} \right| \tag{5.5}$$

The outputs of the static characterization were the pressure-dependent relative resistance curves, the sensitivity S extracted in the low-pressure range, and the coefficient of determination, R^2, used to describe the goodness of the linear fit for sensitivity extraction.

Durability was assessed by applying repeated press-release cycles at 0.5 kPa and monitoring the resistance variation over 1000 cycles using Labtracer. This experiment was performed to evaluate whether the pressure-dependent response was preserved after repeated mechanical stimulation and to compare the stability of the different rGO-R1 coverage conditions. The outputs of the cyclic characterization were the relative resistance response before and after cycling and the corresponding response loss after 1000 cycles.

Results obtained from the static and cyclic tests are reported in Section 5.8.1.

5.7.2 Dynamic Mechanical Characterization

Dynamic mechanical characterization was performed on both rectangular and cross-shaped sensors to compare their response under controlled time-varying stimulation. This experiment provided an intermediate validation step before physiological recordings by allowing the acquired voltage waveform to be compared with a known mechanical input. The corresponding results are reported in Section 5.8.2.

The setup, illustrated in Figure 5.6b, was based on a vibration test system comprising a vibration exciter and its amplifier, driven by a waveform generator. A 3D-printed tip mounted on the shaker stinger applied the oscillatory stimulus to the sensing diaphragm, with different tips used for the rectangular and cross-shaped geometries. Each sensor was mounted in a dedicated 3D-printed support fixed to a metal stand on the workbench. The imposed motion was independently measured using a laser Doppler vibrometer, and the measured velocity was converted into displacement. The sensor output and reference displacement were acquired through a data acquisition system.

Sinusoidal stimuli were applied at 1 Hz, with displacement amplitudes in the order of hundreds of micrometers. A static displacement offset reproduced the preload normally present during manual pulse-wave acquisition. The stimulus amplitude was progressively increased to evaluate the voltage–displacement response and identify the onset of the output plateau.

The output waveform was evaluated qualitatively against the reference displacement, while signal quality was quantified through post-processed SNR. To estimate the noise component, the frequency band containing the imposed periodic response was removed from the acquired signal, and the RMS value of the residual was calculated over the complete analysis window. To estimate the signal component, complementary filtering was used to retain the periodic response and reject the out-of-band noise; the RMS amplitude was calculated for each extracted cycle and then averaged across the acquisition. The SNR was expressed as:

$$\mathrm{SNR}_{\mathrm{dB}} = 20\log_{10}\left(\frac{V_{\mathrm{sig,rms}}}{V_{\mathrm{noise,rms}}}\right), \tag{5.6}$$

where $V_{\mathrm{sig,rms}}$ is the cycle-averaged RMS amplitude of the filtered response and $V_{\mathrm{noise,rms}}$ is the residual noise RMS amplitude.

5.7.3 Preliminary Pulse-Wave Acquisition

A preliminary physiological assessment was performed to determine whether the complete platform could record arterial pulse-related pressure waveforms from superficial anatomical sites. This was an exploratory proof-of-concept study rather than a structured validation protocol; accordingly, the recordings were not organized as a formal cohort study or analysed through repeated-measurement statistics.

Exploratory recordings were acquired from 3 healthy adult members of the research team solely to verify the operation of the complete acquisition chain under physiological conditions. These measurements were demonstrative rather than a formal human-subject study: no statistical analysis was performed and no population-level conclusions were drawn.

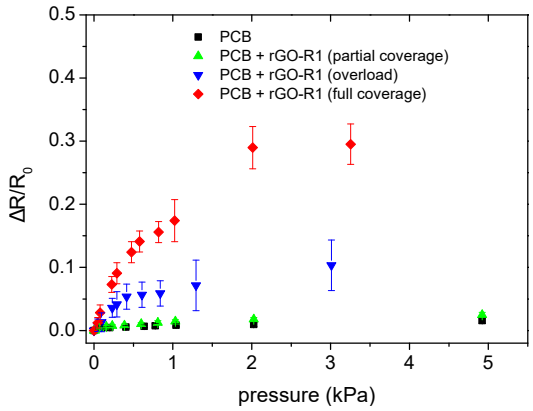
The pen-shaped probes described in Section 5.4 were manually positioned over the carotid and femoral arterial sites. Before each acquisition, the electronic system was initialized, the current-based Wheatstone bridge was calibrated, and the programmable gain was adjusted according to the observed signal amplitude. These operations were performed through the host-software UI, which displayed the two probe channels in real time and supported the adjustment of probe position and contact pressure. The acquired data were stored on the host computer for offline processing.

The experiments included single-site recordings, used to inspect waveform quality and within-session stability, and simultaneous dual-site recordings, used to verify synchronous acquisition from anatomically separated sites. Figure 5.6c shows a representative single-site pulse-wave acquisition on a volunteer.

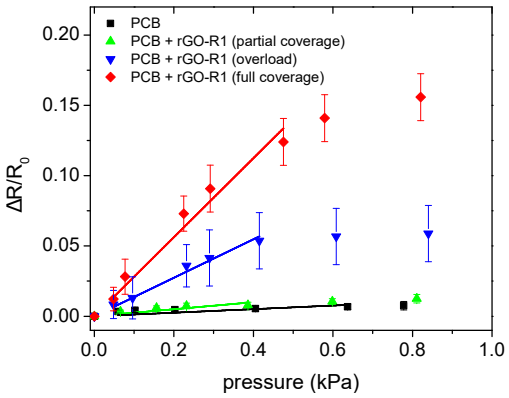
The evaluation focused on qualitative waveform detectability, morphology preservation, probe usability, and subject comfort. A trained healthcare operator also assessed whether the acquired traces were compatible with arterial pulse-wave morphology. The corresponding results are described in Section 5.8.3.

5.8 Results

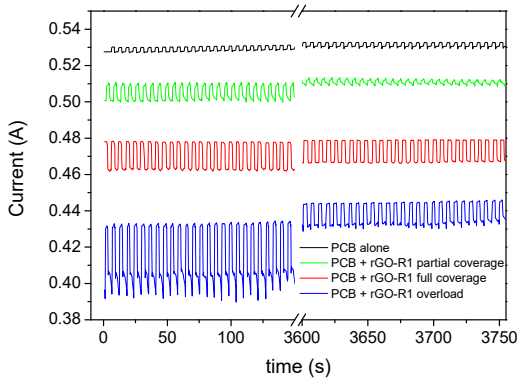
This section reports the experimental results obtained during the characterization and preliminary validation of the rGO-based pressure sensing platform. The results follow



(a) Pressure-dependent resistance response



(b) Low-pressure sensitivity



(c) Cyclic response

Fig. 5.7 Sensor-level response of the rectangular flexible-PCB devices. The plots report relative resistance variation versus pressure (a), sensitivity in the low-pressure range (b), and the response during 1000 cycles at 0.5 kPa (c). Sensitivity error bars denote the standard error of the mean across at least five samples. The solid lines in (b) represent linear fits used to estimate sensitivity in the low-pressure regime.

the validation strategy described in Section 5.7: first, the sensor-level piezoresistive response and cyclic stability are reported; then, the dynamic response of the integrated sensor-readout platform is evaluated under controlled mechanical stimulation; finally, preliminary physiological pulse-wave recordings are presented.

5.8.1 Piezoresistive Response and Cyclic Stability

The sensor-level characterization was used to quantify the pressure sensitivity and fatigue resistance of the rGO-R1-coated flexible PCB sensors.

Table 5.2 Static sensitivity values calculated from the low-pressure response of the flexible PCB pressure sensors.

Sample	Sensitivity (kPa^{-1})	R^2
PCB only	0.013	0.93
PCB + rGO-R1, partial coverage	0.025	0.91
PCB + rGO-R1, full coverage	0.281	0.98
PCB + rGO-R1, overload	0.137	0.99

Figure 5.7a reports the pressure-dependent relative resistance variation, while Figure 5.7b shows the sensitivity extracted in the low-pressure range. These curves show that the amount and distribution of deposited rGO-R1 strongly affected the piezoresistive behaviour. Sensors with partial coverage exhibited the lowest sensitivity among the rGO-R1-coated devices, equal to $0.025\,kPa^{-1}$. This result indicates that an insufficient amount of active material limits the formation of an effective pressure-sensitive conductive path. The highest sensitivity was obtained for fully covered sensors, which reached $0.281\,kPa^{-1}$, with a coefficient of determination of $R^2 = 0.98$ in the low-pressure range. In contrast, overloaded sensors showed a lower sensitivity of $0.137\,kPa^{-1}$, despite the larger amount of deposited active material. This suggests that excessive rGO-R1 coverage does not further improve the piezoresistive response and may reduce the relative contribution of local deformation to the overall conductive path. Sensitivity values and fitting coefficients are summarized in Table 5.2.

Cyclic tests were performed by applying 1000 press-release cycles at 0.5 kPa. As shown in Figure 5.7c, the response reduction after cycling depended strongly on the rGO-R1 coverage condition. Partially covered sensors showed a response loss of approximately 62%, while overloaded sensors showed a comparable loss of approximately 61%. In both cases, the degradation indicates that insufficient or excessive coating negatively affects repeatability under repeated deformation. Fully covered sensors showed better stability, with a response loss of approximately 25% after 1000 cycles. Although a reduction was still present, this condition preserved a larger fraction of the initial response and maintained a more regular dynamic behaviour during repeated loading. The cyclic response was reported using the same relative resistance variation, $\Delta R/R_i$, defined in Equation 5.4. The values measured before and after 1000 cycles, together with the corresponding response losses, are summarized in Table 5.3.

Table 5.3 Relative resistance response at 0.5 kPa before and after 1000 press-release cycles.

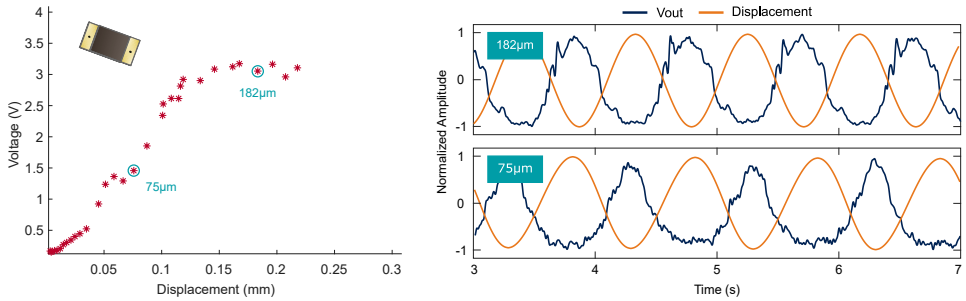
Sample	$\Delta R/R_i$ Initial	$\Delta R/R_i$ After 1000 cycles	Response **loss**
PCB only	0.0038 ± 0.0003	0.0038 ± 0.0002	N.A.
PCB + rGO-R1, partial coverage	0.0157 ± 0.0071	0.0059 ± 0.0011	62%
PCB + rGO-R1, full coverage	0.0335 ± 0.0023	0.0250 ± 0.0031	25%
PCB + rGO-R1, overload	0.0647 ± 0.0040	0.0250 ± 0.0039	61%

Overall, the sensor-level results show that the rGO-R1 coating must be controlled within an appropriate deposition window. Partial coverage does not provide a sufficiently effective active layer, whereas overload reduces the benefit of pressure-induced modulation and degrades stability. Full coverage provided the best compromise between sensitivity and cyclic robustness and was therefore selected as the reference condition for the following validation steps.

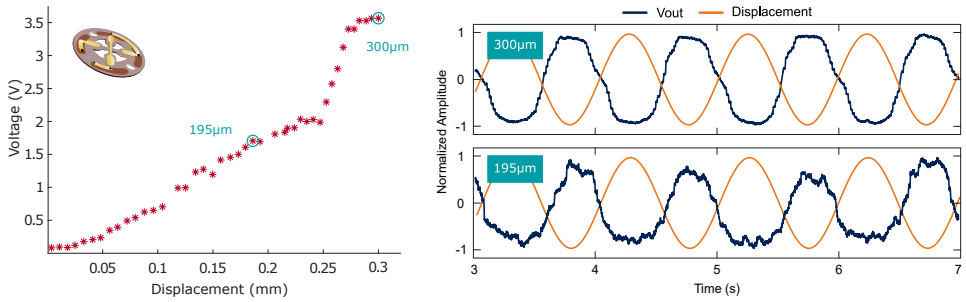
5.8.2 Dynamic Mechanical Response

As described in Section 5.7.2, dynamic mechanical characterization evaluated the response of the integrated sensor-readout platform under controlled sinusoidal stimulation. This test was performed to assess how the output voltage varied with the imposed displacement and whether the acquired waveform preserved the morphology of a known periodic mechanical input.

Figure 5.8 reports the dynamic response of the rectangular and cross-shaped sensors. For each geometry, the voltage–displacement curve was first analysed to identify the useful operating range, the linear response region, and the onset of saturation. The time-domain voltage signal was then inspected at different displacement amplitudes and compared with the corresponding displacement waveform. The figure shows two representative displacement amplitudes for each sensor. Since the operating ranges differed between the two geometries, different amplitudes were selected for representation. For each sensor, the higher amplitude was chosen in the plateau region of the voltage–displacement curve, where the derivative was close to zero. The lower amplitude was selected approximately 100 µm below the plateau, to compare the waveform morphology within a more linear operating region.



(a) Rectangular sensor



(b) Cross-shaped sensor

Fig. 5.8 Dynamic response of the rectangular and cross-shaped sensors. For both rectangular (a) and cross-shaped (b) geometries, the left plot reports output voltage versus measured displacement, while the right plots compare normalized sensor output and reference displacement at representative amplitudes below and at the response plateau. The output voltage has opposite polarity with respect to the imposed displacement because increasing deformation decreases the rGO-R1 sensor resistance through the implemented readout chain.

The rectangular sensor showed higher sensitivity at low displacement amplitudes, up to about 0.10 mm to 0.12 mm. For larger displacements, the curve progressively flattened, reaching a plateau at approximately 182 µm. This behaviour indicates that the rectangular geometry is effective for small deformations, but its useful dynamic range is limited by early saturation. The time-domain waveforms confirmed this behaviour. At 75 µm, the rectangular sensor produced a quasi-sinusoidal voltage output consistent with the imposed displacement. At 182 µm, corresponding to the plateau region, waveform distortion became visible, especially near the peaks. The corresponding SNR values were approximately 21 dB at 75 µm and 20 dB at 182 µm, showing that the increase in displacement did not improve the signal quality once the response approached saturation.

The cross-shaped sensor showed a different behaviour. Its sensitivity at lower displacement amplitudes was lower than the rectangular sensor, but the voltage-

Table 5.4 Representative dynamic mechanical characterization results for rectangular and cross-shaped sensors.

Geometry	Displacement (μm)	SNR (dB)	Observed response
Rectangular	75	21	Quasi-sinusoidal response at low amplitude
Rectangular	182	20	Peak distortion at higher amplitude
Cross-shaped	195	19	Periodic response preserved before plateau
Cross-shaped	300	25	Higher SNR and preserved morphology

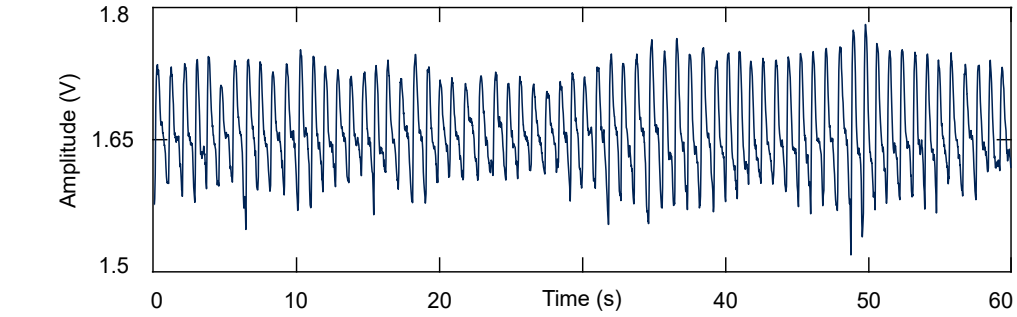
displacement curve remained approximately linear over a wider range, up to about 0.25 mm. After this region, the response increased over a short displacement interval of approximately 0.025 mm, before reaching a plateau at approximately 300 µm. The waveform analysis further supported this interpretation. At 195 µm, the cross-shaped sensor preserved the periodic structure of the imposed displacement, with an SNR of approximately 19 dB. At 300 µm, the SNR increased to approximately 25 dB.

More importantly, for most of the acquisitions performed, the waveform morphology remained closer to the imposed periodic displacement, with less evident distortion than in the rectangular configuration at higher amplitudes. These results are summarized in Table 5.4.

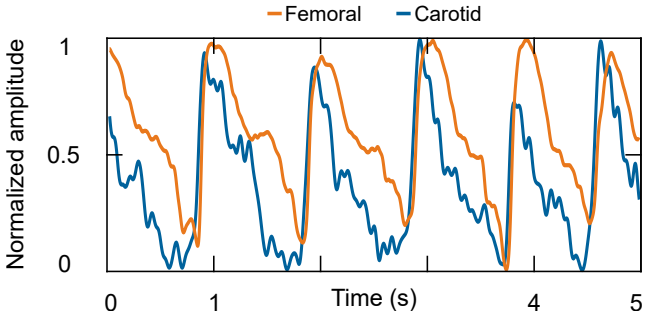
These outcomes support the use of the cross-shaped geometry for the application-oriented sensor. Although the rectangular sensor was more sensitive at lower displacements, it reached saturation earlier and showed greater waveform distortion at higher amplitudes. The cross-shaped sensor provided a broader useful displacement range and better preservation of the dynamic waveform morphology, making it more suitable for pulse-wave acquisition with the final mechanical holder.

5.8.3 Preliminary Pulse-Wave Acquisition

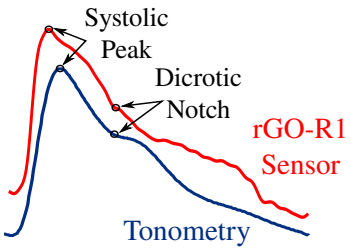
The complete pressure sensing platform was finally evaluated in preliminary physiological acquisitions. As explained in Section 5.7.3, the objective was to verify whether the integrated system could record arterial pulse-related pressure waveforms from superficial anatomical sites using the pen-shaped probes, electronic acqui-



(a) One-minute femoral acquisition



(b) Simultaneous carotid–femoral acquisition



(c) Qualitative morphology comparison

Fig. 5.9 Representative physiological recordings. (a) One-minute femoral acquisition. (b) Five-second simultaneous carotid and femoral acquisition using two probes connected to the same main board. (c) Qualitative comparison of pulse-wave morphology between the rGO-R1 sensor and SphygmoCor applanation tonometry, using separately acquired representative waveforms.

sition system, firmware, and host-software UI described in the previous sections. Representative recordings are reported in Figure 5.9.

Single-site acquisitions showed that the platform could record periodic pulse-related pressure waveforms from superficial arterial sites under manual probe positioning. In particular, prolonged recordings from the femoral artery showed repeated systolic upstrokes and consistent beat-to-beat morphology over the acquisition interval, without evident degradation of the waveform. The representative 1 min femoral recording shown in Figure 5.9a supports the ability of the system to maintain a stable signal within a continuous acquisition session.

The system was also tested in a dual-site configuration using two pen-shaped probes connected to the same main board. This configuration was used to verify whether the platform could acquire pulse-related waveforms from anatomically

separated sites during the same measurement session. As shown in Figure 5.9b, carotid and femoral signals could be recorded simultaneously, and as expected, the femoral waveform exhibited a visible delay with respect to the carotid waveform.

A qualitative comparison with SphygmoCor, the reference device for PWA, was also performed to inspect the morphology of the acquired pulse signal against the gold standard [54, 68]. The two traces were acquired separately, since simultaneous placement of both probes at the same arterial site was not feasible. As shown in Figure 5.9c, the waveform acquired with the developed platform showed the main morphological features expected in an arterial pulse recording, including a rapid systolic upstroke, a dicrotic notch, and a subsequent decay phase. This comparison provided a qualitative indication of morphological compatibility with the reference technique and should not be interpreted as a quantitative agreement analysis.

From a usability perspective, the pen-shaped probe allowed the operator to position the sensing head on both carotid and femoral sites simultaneously while observing the acquired waveform in real time through the host-software UI. This facilitated the adjustment of probe placement and contact pressure during acquisition. No major discomfort was reported during these preliminary sessions, although the assessment of comfort was qualitative and limited to the restricted volunteer group involved in the proof-of-concept tests.

Overall, these preliminary physiological acquisitions demonstrate the feasibility of using the rGO-based pressure sensing platform for non-invasive arterial pulse-wave recording. The results confirm that the sensor, mechanical interface, electronic readout, firmware, and host software can operate together in a realistic measurement scenario, including relatively prolonged single-site recordings and simultaneous dual-site acquisition. At the same time, the acquisitions should be interpreted as proof-of-concept results: further validation on a structured protocol, with controlled positioning, reference instrumentation, and a larger subject cohort, would be required to assess clinical performance.

5.9 Discussion

The results presented in this chapter demonstrate the feasibility of a custom graphene-based pressure sensing platform for non-invasive arterial pulse-wave acquisition. The

platform integrates an rGO-R1 piezoresistive layer, a flexible-PCB sensor structure, a dedicated mechanical holder, current-based Wheatstone bridge electronics, embedded firmware, and host software. This system-level approach is relevant because pulse-wave acquisition depends not only on material sensitivity, but also on mechanical coupling, contact pressure, electronic conditioning, and real-time acquisition control.

At the sensor level, the rGO-R1-coated flexible-PCB devices exhibited a clear pressure-dependent response. Full coverage provided the highest sensitivity among the tested conditions, reaching $0.281\,\mathrm{kPa}^{-1}$ in the low-pressure range, and the best relative stability after repeated loading. Compared with the previous rectangular ITO–PET implementation, which reported $0.82\,\mathrm{kPa}^{-1}$ [94], the flexible-PCB platform showed lower sensitivity but enabled reproducible and scalable manufacturing.

The geometry optimization further showed that sensitivity alone did not determine application suitability. The rectangular sensor had a steeper response at low displacement but reached saturation earlier and exhibited greater waveform distortion. The cross-shaped geometry provided a wider useful displacement range and better waveform preservation during controlled stimulation, making it more appropriate for the pen-shaped probe. The controlled dynamic tests evaluated waveform response at 1 Hz; the broadband frequency response and possible mechanical resonances, including the contribution of the recovery spring, were not quantitatively characterized.

The physiological acquisitions demonstrated proof-of-concept operation at superficial arterial sites. The representative femoral, carotid–femoral, and morphology-comparison recordings showed recognizable arterial pulse morphology, within-session signal stability, and simultaneous dual-site acquisition. The comparison with SphygmoCor was qualitative and should not be interpreted as a quantitative agreement or clinical validation.

The experimental scope and the mechanical robustness of the sensors remain important limitations. Physiological testing involved a restricted volunteer group and was not organized as a structured validation protocol with repeated-measurement statistics or quantitative comparison against reference instrumentation. Probe positioning and preload were manually controlled, as in applanation tonometry, but the flexible sensing diaphragm was particularly sensitive to excessive or off-axis loading. The controlled 1000-cycle tests applied repeatable axial pressure and therefore did not reproduce all loading conditions encountered during manual

acquisition. After some physiological sessions, cracks or local damage were observed in several sensors, and their subsequent response became less predictable. The recordings reported in this chapter were obtained from sensors that were operating correctly; damaged sensors were excluded from further use. Although the damage did not materially affect the acquisitions presented here, it limited sensor reuse and was one of the main reasons for not extending the physiological campaign or pursuing this configuration as a wearable device.

A further limitation arose from the low sensor baseline resistance and the $30\,\Omega$ bridge resistors required to preserve readout sensitivity. Biasing the bridge between V_{CC} and ground increased current consumption and raised the possibility of local heating at the sensor and resistor levels. Since temperature was not measured systematically, its contribution to drift or long-term degradation could not be quantified. These electrical and mechanical constraints represent significant obstacles for wearable operation, where power consumption, thermal comfort, repeatability, and long-term stability are stricter requirements.

5.10 Conclusions

This chapter presented a custom graphene-based pressure sensing platform for non-invasive arterial pulse-wave acquisition. The system combined an rGO-R1 flexible pressure sensor, a pen-shaped mechanical interface, current-based Wheatstone bridge readout electronics, embedded firmware, and host software implementing a GUI for real-time visualization and data storage. In this form, the platform provided a complete acquisition chain, from pressure transduction at the skin interface to digital recording of pulse-related waveforms.

The platform demonstrated the potential of integrating a custom graphene-based sensor with dedicated mechanics and electronics for arterial pulse-wave acquisition. Its main strengths were the application-oriented design, the possibility of investigating the full chain from material-level piezoresistive response to physiological recording, and the ability to acquire representative single-site and dual-site pulse waveforms. At the same time, its mechanical fragility under uncontrolled pressure, power-consumption and potential self-heating concerns, and limited physiological validation restricted its direct translation into a robust wearable PWV-monitoring system.

The following chapter moves toward a wearable platform based on commercial piezoresistive sensors, prioritizing robustness, repeatable body integration, dual-site acquisition stability, and PWV-oriented monitoring.

Chapter 6

Wearable Device for Pulse Wave Velocity Monitoring

6.1 Introduction and Motivation

The previous chapter presented a custom rGO-R1-based pressure sensing platform and demonstrated the feasibility of acquiring arterial pulse-related waveforms from superficial anatomical sites using a flexible piezoresistive sensor, dedicated mechanical coupling, and a custom electronic readout chain. The experience gained with sensor deformation, elastic recovery, and electronic conditioning informed the subsequent development of a wearable platform for pulse wave velocity (PWV) assessment. At this stage, the focus shifted from custom sensor development and preliminary waveform acquisition towards repeated use, wearable integration, and clinical evaluation. The practical constraints discussed in Sections 5.9 and 5.10 therefore motivated the adoption of commercial sensing elements.

Alongside sensor integration, a further design objective was to reduce the need for continuous manual probe stabilization. Variations in anatomical localization, contact pressure, probe orientation, and stability can introduce operator-dependent variability into the acquired waveforms and the resulting PWV estimate [103]. Even when measurements are performed by trained personnel, maintaining comparable coupling conditions across repeated acquisitions remains difficult.

Building on this experience, the work presented in this chapter, also reported in [4], developed a wearable and wireless system based on commercial force-sensing resistors (FSRs) for one-step carotid–femoral PWV assessment in supervised hospital settings. Commercial FSRs provide thin, flexible, low-cost, replaceable, and readily available sensing elements suitable for integration into compact body-worn modules. Their availability and documented characteristics allowed development to focus on mechanical fixation, electronic conditioning, and synchronized acquisition.

The system is intended for supervised use by trained personnel. The operator must identify the arterial sites, position the modules, adjust the fixation, and verify signal quality before acquisition. After positioning, elastic fixation supports the modules without continuous manual stabilization. Battery power and wireless communication minimize cabling around the subject.

The system adopts a one-step acquisition strategy, in which the carotid and femoral waveforms are recorded simultaneously, as described in Section 4.3.3. Each of the two sensing modules integrates an FSR, analog signal conditioning, embedded acquisition, battery power, and Bluetooth Low Energy (BLE) communication. Before placement, a direct hardware interface establishes a common timing reference between the modules. During acquisition, a receiver connected to the host computer collects the transmitted signals. The host software provides acquisition control, real-time visualization, and data storage, while offline processing identifies fiducial points, estimates pulse transit time (PTT), and calculates PWV.

Sections 6.2–6.6 describe the sensing modules, wearable integration, electronics, firmware, synchronization, and host software. Section 6.7 presents the PTT and PWV estimation workflow, while Sections 6.8 and 6.9 report the validation methodology and results. The chapter closes with a discussion of the system-level performance, limitations, and future development directions.

6.2 Design Strategy and System Overview

The system was designed around the requirements of one-step carotid–femoral PWV assessment. It therefore had to preserve the relative timing of signals recorded at two anatomically separated sites, remain fixed to the body during acquisition,

operate without wired connections to the host computer, and provide data suitable for quantitative PWV estimation.

6.2.1 Requirements for Wearable cfPWV Monitoring

The waveform-preservation and bandwidth considerations introduced in Section 5.2.1 also apply to this platform. The following requirements reflect its intended use in short, supervised cfPWV measurement sessions:

- Dual-site timing preservation: The carotid and femoral waveforms must be acquired simultaneously with a common timing reference and a temporal grid of 1 ms or finer for fiducial-point localization. This target is motivated by the high-resolution foot-detection guidance provided for PWV reference measurements [57].
- Stable wearable coupling: The sensors must remain positioned over the target arteries throughout each recording without continuous manual stabilization. Fixation must allow adjustment of contact preload and support comfortable acquisition.
- Compact battery-powered operation: Each module must integrate sensing, conditioning, digitization, wireless communication, and battery power within a wearable enclosure. For the intended session-based use, an indicative autonomy target is at least 2 h of active acquisition and wireless transmission between charges, allowing repeated short recordings and setup time.
- Data handling and PTT extraction: Both waveform streams must be available for real-time inspection and stored throughout each recording, preserving channel identity and relative timing for offline fiducial-point extraction and PTT estimation.

6.2.2 System Overview

An overview of the complete system architecture is shown in Figure 6.1. The system comprises four functional units: two wearable sensing modules, one for each anatomical site, a BLE receiver, and a host computer. The sensing modules

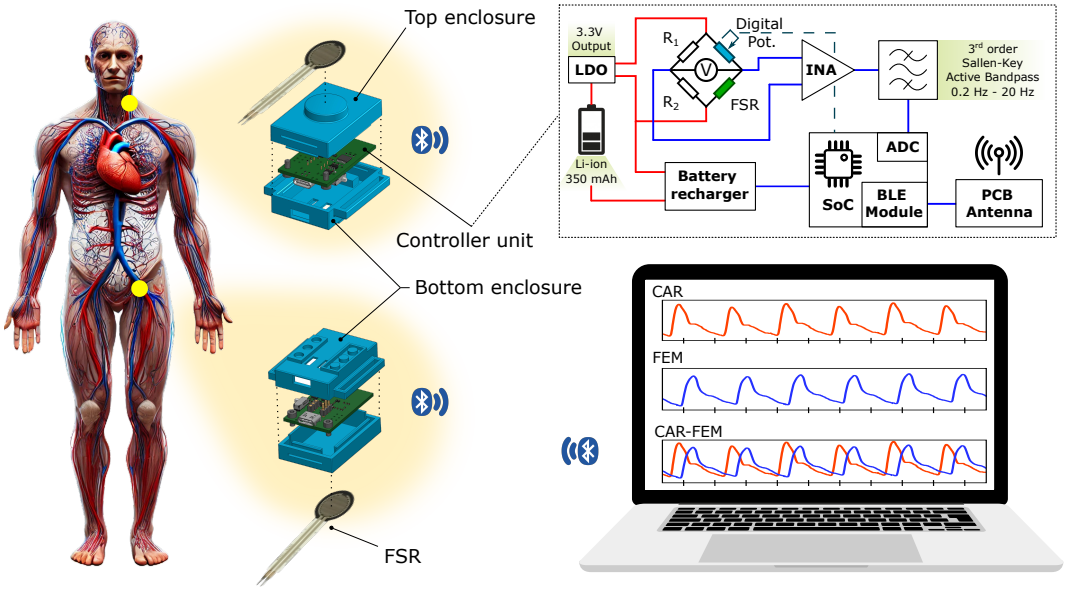


Fig. 6.1 System architecture of the wearable PWV platform. Two body-worn analog front-end (AFE) modules, positioned over the carotid and femoral arteries, integrate an FSR, resistance-to-voltage conversion, analog conditioning, 12-bit digitization, and BLE transmission. Before placement, the modules are temporarily coupled through a hardware interface to establish a common timing reference. During acquisition, the synchronized data streams are transmitted to a dedicated receiver and forwarded to a host computer, where a graphical user interface (GUI) supports real-time visualization, data storage, and subsequent PTT and PWV estimation. Adapted from [4].

are electrically equivalent and are identified in software as the carotid and femoral units according to their anatomical placement. Each module is battery-powered and integrates an FSR, analog conditioning, an embedded unit, and BLE communication.

Local force variations produced by arterial pulsation change the FSR resistance. These variations are converted into conditioned voltage waveforms, digitized, and transmitted wirelessly to the receiver connected to the host computer.

Before positioning, the sensing units are temporarily coupled through their hardware synchronization interface, which establishes a common timing reference between the modules. They are then separated, positioned over the carotid and femoral sites, and used for synchronized acquisition. The two waveforms are displayed in real time and stored for subsequent processing. Corresponding fiducial points are identified offline to estimate PTT, which is combined with the measured anatomical distance to calculate PWV.

Table 6.1 Main specifications of the wearable PWV monitoring system [4].

Specification	Value
Sensor	Interlink FSR 402
Analog passband	0.2 Hz to 20 Hz
MCU	STM32WB55 SoC
ADC resolution	12-bit
Sampling frequency	1 kHz per signal
Wireless interface	BLE 5.0
BLE payload per node	16 bytes; 8×16-bit samples
BLE packet rate per node	125 packets/s
Useful payload rate per node	$2\,\mathrm{kB\,s^{-1}}$
Receiver output block	32 bytes; $4\,\mathrm{kB\,s^{-1}}$ aggregate
Battery	3.7 V; 350 mA h
Estimated active current	Approximately 6 mA
Enclosure dimensions	45 mm×42 mm×20 mm

Table 6.1 summarizes the principal specifications of the proposed system.

6.3 Sensor Selection and Working Principle

FSRs are thin-film piezoresistive transducers whose electrical resistance decreases as the force applied to their active area increases. They typically consist of flexible polymer layers incorporating conductive electrodes and pressure-sensitive resistive ink. When subjected to an external force, deformation of the layered structure modifies the electrical contact within the sensing region, producing a force-dependent resistance variation.

FSRs are particularly suitable for the present application due to their low thickness, mechanical flexibility, environmental robustness, low cost, and ease of integration. These characteristics allow the sensing element to detect small variations in force associated with arterial pulsation while remaining compatible with a wearable module.

Commercial thin-film piezoresistive force sensors are generally implemented in either *shunt-mode* or *thru-mode* configurations. In thru-mode devices, current flows through the thickness of the pressure-sensitive layer between opposing electrodes. In shunt-mode devices, such as the Interlink FSR family, pressure increases the

contact between a resistive layer and interdigitated electrodes, creating lateral conductive paths and reducing the measured resistance [104, 105]. Comparative studies by Lebossé *et al.* evaluated the static and dynamic behaviour of shunt-mode Interlink FSRs and thru-mode Tekscan FlexiForce sensors [105, 106]. Although the thru-mode sensor exhibited better static linearity, it also showed more pronounced nonlinearities under dynamic loading. The shunt-mode Interlink FSR provided a more favorable dynamic response, making it more suitable for the time-varying mechanical stimulation associated with arterial pulse waves.

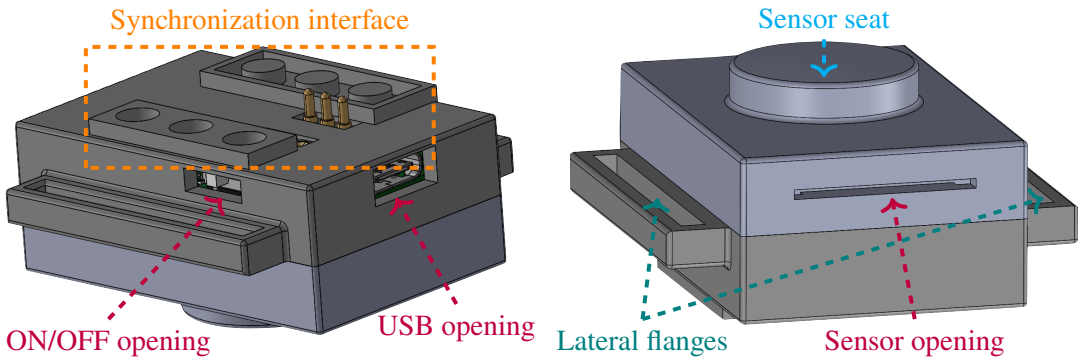
For this reason, the shunt-mode Interlink Electronics FSR 402 was selected in the proposed system. According to the manufacturer documentation for the specific sensor revision, the FSR 402 has an approximate force-sensing range of 0.1 N to 10 N and a stated durability of up to 10^7 actuations. The reported ±2% repeatability refers to a single device under controlled and repeatable actuation and should not be interpreted as the repeatability of the complete wearable measurement [107, 104].

During pulse-wave acquisition, the sensor is subjected to a static preload that maintains contact with the skin and establishes its operating point. Arterial pulsation produces a smaller dynamic force variation superimposed on this preload, periodically modifying the sensor resistance.

6.4 Wearable Integration and Sensor Coupling

The wearable integration strategy was developed to maintain stable coupling between the FSR and the body while preserving portability, user comfort, and access to the carotid and femoral measurement sites. The mechanical design had to accommodate and protect the sensing and electronic components, maintain the sensor in contact with the skin without continuous manual stabilization, and allow the position and preload to be adjusted according to the anatomical site and the individual subject.

The system consists of two identical body-worn modules, each integrating the FSR, acquisition electronics, wireless communication circuitry, and battery within a dedicated enclosure. The housing was designed and manufactured by 3D printing. The rendered views in Figures 6.2a and 6.2b illustrate the enclosure design, while Figure 6.2c shows a photograph of the assembled modules. Lateral flanges allow an elastic band to be inserted for wearable fixation. The non-active area of the FSR



(a) External view

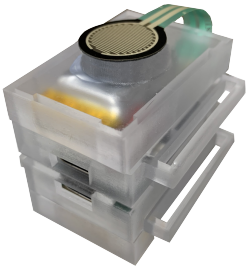
(b) Sensor-side view



(c) Photograph of the modules



(d) Wearable system during clinical testing



(e) Enclosure matching for synchronization

Fig. 6.2 Wearable integration of the PWV monitoring system [4]. (a) Render of the external side of the enclosure, including the region for the hardware synchronization interface and openings for ON/OFF switch and USB cable. (b) Render of the sensor-side view, showing the FSR seat at the skin-facing region, opening for the sensor to reach the readout electronics, and lateral flanges for the elastic band. (c) Photograph of the assembled wearable modules. (d) Example of carotid and femoral module positioning on a participant during the clinical validation study. (e) Coupled enclosures for module synchronization.

substrate extends outside the enclosure and is bent over a raised seat on the skin-side of the housing, where the active area of the FSR is fixed with glue.

The two modules can also be coupled through an interlocking enclosure geometry. When joined, spring-loaded contacts contact the corresponding mating pads on the opposite module, providing the connections required for hardware synchronization, described in Section 6.5.2.

For measurement, the modules are positioned over the carotid and femoral arteries and secured using the elastic bands, as shown in Figure 6.2d. A trained operator identifies each arterial site by palpation, aligns the active area of the FSR with the point where pulsation is most evident, and adjusts the band tension to obtain stable

sensor–skin contact without excessive compression. The operator then verifies the acquired waveforms through the graphical user interface and adjusts the position or tension when necessary. Once suitable coupling has been achieved, the modules remain fixed throughout the recording without continuous manual handling. During the validation study, no participant reported discomfort or skin irritation associated with the wearable system [4].

6.5 Electronic Acquisition System

The electronic acquisition system of each wearable module is implemented on a single PCB and comprises two main functional blocks: an analog front-end and an embedded unit. The analog front-end converts the force-dependent resistance variations of the FSR into a conditioned voltage waveform, while the embedded unit performs digitization, control, wireless communication, hardware synchronization, and power management. The two PCBs are electrically identical and can therefore be used interchangeably at the carotid and femoral sites.

The integration of analog, digital, power, and radio-frequency circuitry within a compact wearable module required careful functional partitioning. In particular, the small pulse-related component had to be preserved in the presence of digital switching activity and BLE transmission.

6.5.1 Analog Front-End

The analog front-end converts the small pulse-induced resistance variations of the FSR, superimposed on a placement-dependent static baseline, into a band-limited voltage signal suitable for digitization. The acquisition chain consists of a Wheatstone bridge with digital balancing, an instrumentation-amplifier stage, and analog high-pass and low-pass filters.

The FSR is inserted into one arm of a voltage-output Wheatstone bridge, supplied between 3 V and ground. Force-induced changes in its resistance unbalance the bridge and generate a differential voltage. In contrast to the current-based architecture used for the low-resistance rGO-R1 sensor (Section 5.5.1), the higher operating resistance

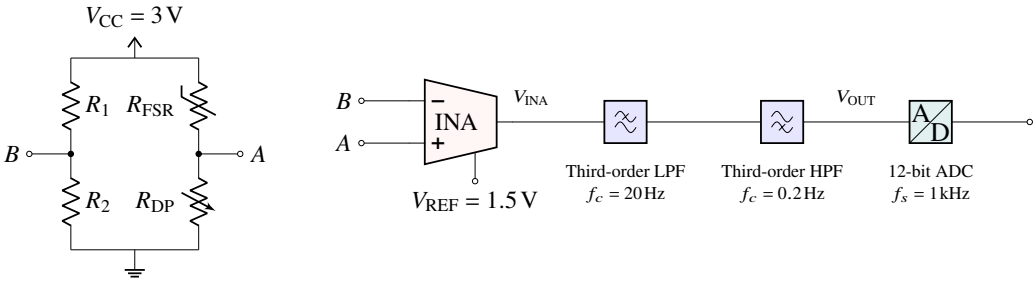


Fig. 6.3 Simplified analog front-end of each wearable sensing module. The FSR and the SPI-controlled digital potentiometer form one branch of a 3 V Wheatstone bridge, while two fixed resistors form the opposite branch. The differential bridge output is processed by an instrumentation amplifier referenced to 1.5 V. Third-order low pass and high-pass stages define the approximate 0.2 Hz to 20 Hz analog passband before ADC digitization at 1 kHz.

of the commercial FSR allows it to be interfaced effectively using a conventional voltage-output bridge.

A digital potentiometer is included in the same bridge branch to compensate for the static operating point introduced by sensor placement and elastic-band tension. Its resistance is adjusted under firmware control according to the calibration procedure described in Section 6.6.1.

The differential bridge output is processed by an instrumentation amplifier referenced to 1.5 V. The reference voltage centers the pulse-related variation within the unipolar range of the subsequent analog stages and ADC. The implemented instrumentation-amplifier introduced unity gain.

As summarized in Figure 6.3, the amplifier output is processed by third-order low-pass and high-pass filters with cutoff frequencies of approximately 20 Hz and 0.2 Hz, respectively. Together, these stages attenuate high-frequency interference and slow baseline variations while preserving the pulse-wave components required for fiducial-point extraction. The final conditioned signal is digitized by the ADC integrated into the embedded system-on-chip.

6.5.2 Embedded Unit

The embedded unit is based on the STM32WB55 system-on-chip (SoC), which integrates processing cores, digital peripherals, a 12-bit ADC, and a BLE 5.0 radio [108]. The SoC controls the bridge-balancing potentiometer, digitizes the

conditioned pulse waveform, manages the synchronization interface, and transmits the acquired samples to the external receiver. Integrating acquisition and wireless communication in the same device reduced the PCB area and eliminated the need for a separate radio module.

The ADC samples the conditioned waveform at 1 kHz, substantially oversampling the nominal 0.2 Hz to 20 Hz analog passband. The resulting 1 ms sample interval provides a dense temporal grid for locating carotid and femoral fiducial points and limits sampling-related quantization in PTT estimation.

BLE communication is supported by an antenna printed on the PCB and its associated filtering and impedance-matching network. Because the antenna operates close to the body and within a compact enclosure, its layout follows the manufacturer's geometry and copper keep-out recommendations [109]. The antenna was designed to provide reliable short-range communication between each wearable module and the dedicated receiver rather than long-range wireless coverage.

The PCB also integrates the physical interface used to synchronize the two wearable units prior to placement. When the enclosures are coupled, as shown in Figure 6.2e, spring-loaded contact pins on each module engage the mating contact pads on the opposite module, electrically connecting the synchronization lines. This connection produces a digital transition on a dedicated GPIO input of each SoC. The firmware uses the detected edge to align the acquisition time bases of the two modules. The corresponding synchronization sequence and firmware operation are described in Section 6.6.1.

Each PCB is powered by a 3.7 V, 350 mA h rechargeable lithium-ion battery. A dedicated battery-charging circuit supports recharging through USB-C, while a low-dropout regulator provides a stable 3 V supply for the analog, digital, and radio-frequency circuitry. Local decoupling and supply filtering were included to limit the propagation of digital and radio-frequency disturbances into the analog front-end.

The average current consumption during continuous acquisition and BLE streaming was estimated at approximately 6 mA. This value was derived from a parallel study using the same sampling frequency, ADC format, and BLE communication configuration [89]. It suggests that the selected battery capacity is sufficient for a full day of supervised hospital use. However, it should be interpreted as an

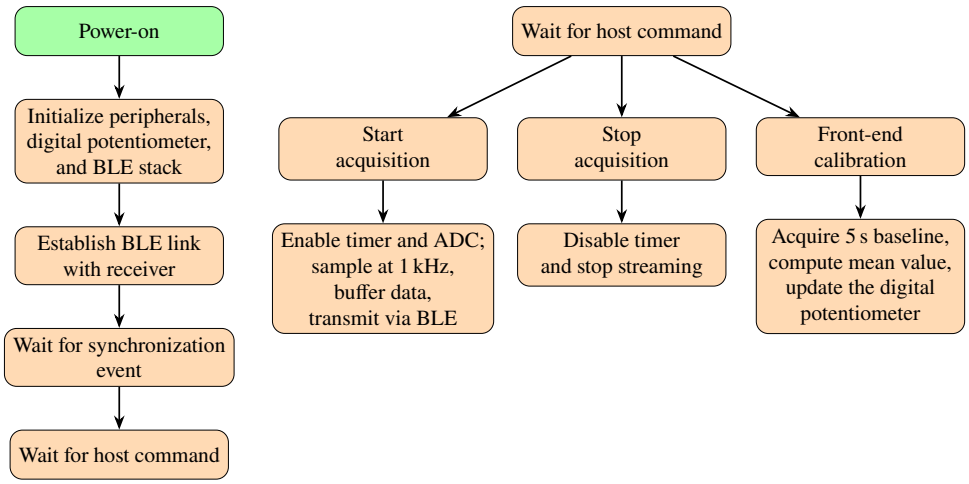


Fig. 6.4 Operating sequence of each wearable sensing node. Module decoupling generates the hardware event used to initialize the local acquisition time base. The host can then trigger front-end calibration or timer-driven sampling. Samples are grouped into eight-sample packets and transmitted to the receiver through BLE.

order-of-magnitude estimate rather than as a complete experimental characterization of autonomy under all operating conditions.

6.6 System Control and Wireless Communication

System operation is coordinated across three software layers: the firmware running on each wearable node, the firmware of the dedicated BLE receiver, and the host software. The node firmware manages synchronization, front-end calibration, signal acquisition, buffering, and wireless transmission. The receiver collects the two BLE data streams and forwards them to the host computer through USB. The graphical user interface provides system-control commands, real-time waveform visualization, and data storage. The same control logic is used by both wearable nodes, which are identified according to the anatomical site assigned during the measurement.

6.6.1 Operating Sequence and Front-End Calibration

The firmware operating sequence is summarized in the workflow diagram in Figure 6.4. At startup, the node firmware initializes the ADC, timers, SPI interface, synchronization GPIO, and BLE communication stack. The digital potentiometer

included in the Wheatstone bridge is initially set to a predefined resistance selected to place the analog front-end close to its nominal operating point.

Before sensor placement, the two modules are temporarily coupled through the hardware synchronization interface described in Section 6.5.2. Decoupling causes the synchronization GPIO of each module to transition from low to high. The corresponding interrupt resets and starts the local timer used to schedule acquisition and firmware operations. The two nodes therefore establish their time bases from the same hardware event, independently of BLE transmission latency. The residual synchronization offset is determined by contact-transition skew, GPIO and interrupt latency, and timer-clock granularity. Even if their contributions were not independently characterized, their expected scale is small compared with the 1 ms sampling interval.

After positioning, an initial 5 s interval is used to calibrate the sensing interface. The firmware calculates the mean value of the acquired signal and adjusts the digital potentiometer through SPI. This compensates for the static pressure introduced by the local anatomy and fixation tension, reduces the bridge offset, and places the pulse-related component within the useful input range of the analog acquisition chain. Once calibration is complete, the system enters the normal acquisition and transmission phase.

6.6.2 BLE Data Transmission and Receiver Interface

BLE was selected to support continuous transmission from the wearable nodes without wired connections to the host computer. The standard BLE 5.0 stack provided for the STM32WB55 platform was used without protocol-level modifications [4].

During acquisition, each wearable node samples one pulse-wave channel at 1 kHz. Each digitized sample is stored in a 16-bit word, comprising the 12-bit ADC value and the associated identification bits. Samples are grouped into BLE packets containing 8 x 16-bit words, corresponding to a payload of 16 bytes. Consequently, each node transmits 125 BLE packets per second, with a useful payload rate of approximately $2\,\mathrm{kB\,s^{-1}}$ per wireless link.

The receiving dongle collects one packet from each node and reorganizes the samples into a single dual-channel block containing 16 interleaved 16-bit samples, eight from the carotid node and eight from the femoral node. Each combined block

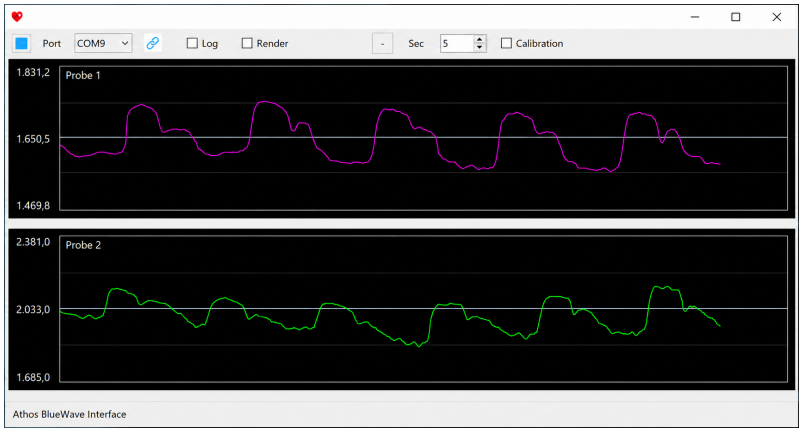


Fig. 6.5 Graphical user interface screenshot.

therefore contains 32 bytes and represents 8 ms of data. The dongle forwards the data to the host computer through USB, with a data rate of $4\,\mathrm{kB\,s^{-1}}$.

The 1 kHz sampling frequency limits the temporal extent of gaps caused by isolated missing samples. Moreover, the beat-based PTT estimation method evaluates multiple fiducial-point pairs within each acquisition window and is therefore tolerant to sporadic missing data [87]. Nevertheless, packet-loss performance was not independently characterized under all operating and environmental conditions.

6.6.3 Host Software and User Interface

The host software was developed in C# using Microsoft Visual Studio. It provides the operator with control of the acquisition workflow and real-time access to the carotid and femoral waveforms. The graphical user interface manages communication with the receiving dongle, issues acquisition commands, displays the two pulse-wave signals, and stores the received data for offline processing. A screenshot of the GUI is displayed in Figure 6.5.

Separate plots are provided for the carotid and femoral channels, allowing the operator to assess waveform amplitude and morphology during module placement. This real-time feedback supports adjustment of the sensor position and fixation tension before the final recording. Data received through the virtual COM interface are separated into channel-specific buffers for visualization and saved in binary format for subsequent PTT and PWV analysis.

6.7 Signal Processing and PWV Estimation

The carotid and femoral waveforms were processed offline to obtain one PWV estimate from each acquisition. The processing workflow comprises three main stages: selection and filtering of the analysis window, extraction and pairing of corresponding fiducial points, and calculation of PTT and PWV [4]. The implemented method combines elements of previously developed pulse-wave processing algorithms and was adapted to the signals acquired by the proposed wearable system [87, 110].

6.7.1 Signal Pre-processing

The acquired waveforms were sampled at 1 kHz. For each recording, the final 12 s interval was selected, and its last 2 s were discarded, resulting in a 10 s analysis window. This duration was selected to ensure comparability with the recordings processed by the clinical reference system. The final 2 s were excluded because operator interaction at the end of the acquisition could introduce motion artifacts.

The selected carotid and femoral signals were digitally filtered using Butterworth filters applied independently and identically to the two channels. A fourth-order low-pass prototype with a cutoff frequency of 10 Hz was followed by a fifth-order high-pass prototype with a cutoff frequency of 0.5 Hz. Both stages were applied in the forward and reverse directions using the MATLAB `filtfilt` function. This procedure provided zero-phase filtering and therefore avoided filter-induced displacement of the fiducial points. Because of the forward–reverse implementation, the effective magnitude-response orders were eight and ten for the low-pass and high-pass stages, respectively.

6.7.2 Fiducial-Point Extraction and PWV Calculation

Several methods have been proposed to identify pulse-wave fiducial points and estimate the temporal delay between arterial signals, including the ITP method and approaches based on the second derivative [111]. The ITP method was adopted because of its widespread use in clinical PWV systems and its robustness under different waveform conditions.

The ITP extraction procedure is illustrated in Figure 6.6. Zero padding was first added at the beginning and end of each signal to facilitate feature detection near the boundaries of the analysis window. A further low-pass-filtered version of the waveform, with a cutoff frequency of 1.5 Hz, was then generated. The minima of this smoothed signal provided robust initial references for locating the foot of each pulse wave, as shown in Figure 6.6a [87].

For each cardiac cycle, the pulse-wave foot was identified in the neighborhood of the corresponding reference minimum, as shown in Figure 6.6b. The point of maximum first derivative was then located along the systolic upstroke, and the tangent to the waveform at this point was calculated, as shown in Figure 6.6c. The intersection between this tangent and the horizontal line passing through the local foot level defined the ITP, as illustrated in Figure 6.6d. The same procedure was applied independently to the carotid and femoral waveforms.

After ITP extraction, carotid and femoral points associated with corresponding cardiac cycles were paired. Fiducial points without a counterpart in the opposite waveform were discarded. For each accepted pair, the temporal difference between the femoral and carotid ITPs provided a beat-level PTT estimate.

Pairs producing a delay outside the interval defined by the mean PTT plus or minus 1.5 times its standard deviation were classified as outliers and excluded. The final PTT for the acquisition was calculated as the mean delay across the remaining pairs:

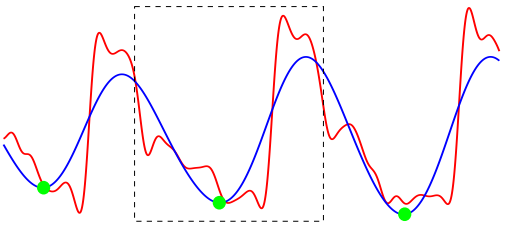
$$PTT = \frac{1}{N} \sum_{i=1}^{N} \left(t_{\mathrm{fem},i} - t_{\mathrm{car},i} \right), \tag{6.1}$$

where $t_{\mathrm{car},i}$ and $t_{\mathrm{fem},i}$ are the carotid and femoral ITP times of the i-th accepted cardiac cycle, respectively, and N is the number of retained pairs.

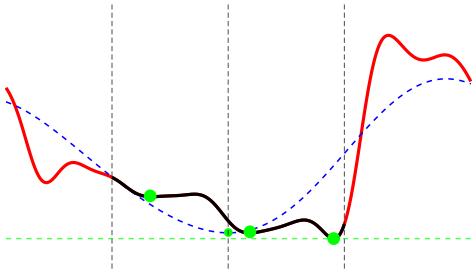
PWV was subsequently calculated by combining the estimated PTT with the measured carotid–femoral distance:

$$PWV = \frac{0.8\, d}{PTT}, \tag{6.2}$$

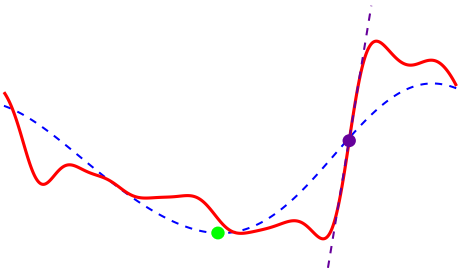
where d is the tape measured carotid–femoral distance and the coefficient 0.8 approximates the effective arterial propagation path from the corresponding surface



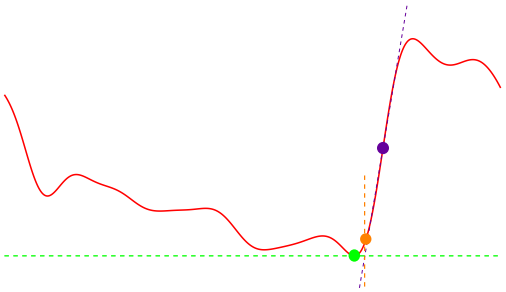
(a) Original waveform (red), 1.5 Hz low-pass-filtered waveform (blue), and detected reference minima (green).



(b) Identification of the local pulse-wave foot (green).



(c) Tangent evaluated at the point of maximum derivative (purple).



(d) Resulting ITP (orange).

Fig. 6.6 Processing steps used to identify the ITP: (a) extraction of reference minima from the 1.5 Hz low-pass-filtered waveform; (b) identification of the local pulse-wave foot; (c) calculation of the tangent at the point of maximum derivative; and (d) identification of the ITP from the intersection between the tangent and the horizontal line passing through the foot level. Numerical time and amplitude scales are omitted because they vary among acquisitions. Adapted from [4].

distance [69]. This calculation follows the cfPWV distance convention introduced in Section 4.3.3 and Equation 4.4.

6.8 Validation Methodology

The complete wearable system was evaluated against SphygmoCor Vx in a supervised clinical study conducted in collaboration with the medical team of the A.O.U. Città della Salute e della Scienza di Torino. The objective was to assess the agreement and repeatability of the proposed system relative to the clinical reference device [4].

6.8.1 Study Design and Population

The study included 101 adult volunteers. The Bioethics Committee of the University of Turin approved the study under Protocol No. 161272 on 8 March 2021. Experienced medical professionals conducted the recordings at the Candiolo Cancer Institute FPO-IRCCS, Division of Internal Medicine, Department of Medical Sciences. All participants provided informed consent for the use of their data for research purposes. Participant recruitment was managed by the clinical team at the Candiolo Cancer Institute FPO-IRCCS.

The study design was informed by recent recommendations for the comparison of devices used for PWV assessment, with the objective of obtaining a heterogeneous cohort while excluding conditions expected to substantially affect arterial stiffness [57]. Participants had no known history of cardiovascular disease. Exclusion criteria included body mass index above $40\,\mathrm{kg\,m^{-2}}$, previous stroke, diabetes, hypertension, pregnancy, or other significant medical conditions affecting cardiovascular function.

The main demographic characteristics of the study population are summarized in Table 6.2.

6.8.2 Measurement Procedure and Reference System

SphygmoCor Vx was adopted as the clinical reference system. Each participant underwent three PWV measurements with the proposed wearable device and three measurements with SphygmoCor. All measurements were performed within the same study session, with each complete acquisition requiring a few minutes, including subject preparation and device positioning.

For the proposed system, the two wearable modules were first synchronized and then separated. The operator identified the carotid and femoral arteries by palpation

Table 6.2 Main characteristics of the validation cohort. Continuous variables are reported as mean ± standard deviation.

Characteristic	**Value**
Number of participants	101
Women / men	55 / 46
Height	166 cm ± 10 cm
Weight	69.2 kg ± 13.8 kg
Age < 30 years	17
Age 30–49 years	14
Age 50–69 years	58
Age ≥ 70 years	12

and positioned the modules using the elastic bands, as shown in Figure 6.2d. After initial placement, the front-end calibration procedure described in Section 6.6.1 was performed to compensate for the static preload introduced by local anatomy and fixation tension.

The carotid and femoral waveforms were then inspected in real time through the graphical user interface. When necessary, the operator adjusted the module position or band tension and repeated the calibration procedure until traces with suitable amplitude, morphology, and stability were obtained. A synchronized carotid–femoral recording was then acquired and stored for offline processing.

During recording, the operator continued to monitor both waveforms and terminated the acquisition after a sufficiently stable segment free from evident artifacts had been obtained. This procedure motivated the analysis-window selection described in Section 6.7.1, in which the final portion of the recording was discarded to limit artifacts associated with operator interaction at the end of the measurement.

PWV calculation also required measurement of the surface distance between the anatomical sites. A flexible measuring tape with a resolution of 1 mm was used to record three distances: from the carotid measurement site to the sternal reference point, from the femoral measurement site to the same reference point, and directly between the carotid and femoral measurement sites. These measurements were collected together to document the acquisition locations and preserve consistency across repeated measurements.

The direct surface distance between the carotid and femoral acquisition sites was used as d in Equation 6.2. The factor 0.8 was applied to approximate the

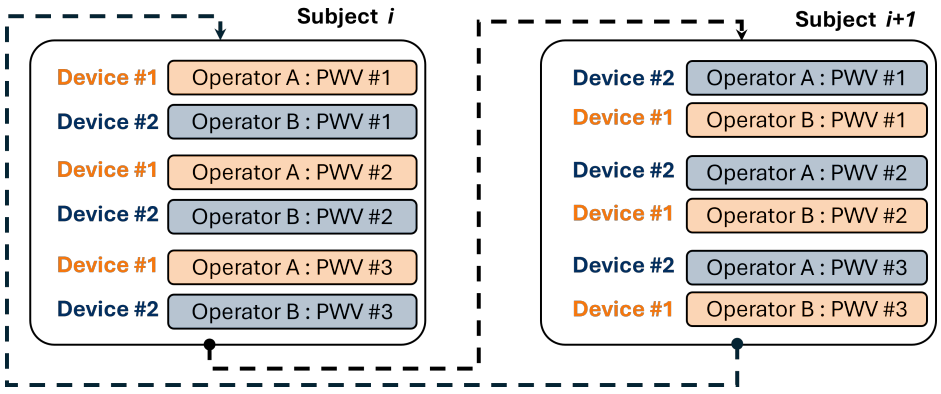


Fig. 6.7 Alternating protocol used for the comparison between the proposed wearable system and SphygmoCor. Measurements with the two devices were alternated to minimize the temporal interval between the methods, while operator assignments were alternated across consecutive participants to reduce operator-dependent bias. Adapted from [4].

effective arterial propagation path from the measured surface distance [69]. The same path-length convention was used for the proposed system and the reference analysis.

Simultaneous measurements were not feasible in the tested configuration because the elastic band securing the wearable carotid module encircled the neck and obstructed access to the contralateral carotid site, preventing placement of the SphygmoCor probe. Acquisitions with the proposed system and SphygmoCor were alternated within each participant to reduce the temporal interval between corresponding PWV estimates. Medical operators were also alternated according to a predefined sequence across consecutive participants. As illustrated in Figure 6.7, the operator who performed the SphygmoCor measurement on participant i performed the proposed-device measurement on participant $i+1$. This strategy was adopted to limit systematic effects associated with device order and operator-specific technique.

Signal quality was assessed both during acquisition and offline. Real-time visual inspection was used to evaluate waveform amplitude, periodicity, morphology, and stability during module placement and recording. An offline quantitative assessment was subsequently performed by calculating the SNR separately for the carotid and femoral waveforms.

The same RMS-based procedure introduced for the rGO-R1 platform in Section 5.7.2 was adopted. Complementary filtering was used to separate the periodic pulse-wave component from the residual noise. The RMS amplitude of the signal

component was calculated for each cardiac cycle and averaged over the analysis window, whereas the RMS amplitude of the residual component was calculated over the complete window. The SNR was then obtained using Equation 5.6. The corresponding values are reported in Section 6.9.1.

For each participant and device, the three PWV estimates were screened before calculating the representative value. A measurement whose absolute relative difference from the median of the corresponding triplet exceeded 15% was classified as an outlier and excluded. The remaining measurements were averaged to obtain one representative PWV estimate per participant and device for the agreement analysis.

6.8.3 Statistical Analysis

Agreement between the proposed wearable system and SphygmoCor was evaluated using the representative PWV value obtained for each participant and device. Pearson's correlation coefficient was used to quantify linear association, whereas Spearman's coefficient was used to assess monotonic association. Lin's concordance correlation coefficient was calculated to evaluate agreement with respect to the identity line, while the intraclass correlation coefficient was used as a complementary measure of agreement between the two methods [112, 113].

A Bland–Altman analysis was performed to quantify systematic and random differences between devices [114]. The mean paired difference represented the measurement bias. To remain consistent with the analysis reported in [4], the limits of agreement were calculated as the mean difference plus or minus two times the standard deviation of the paired differences.

Within-device repeatability was evaluated from the variability of the three repeated measurements obtained for each participant. The mean within-subject standard deviation was calculated separately for the proposed device and SphygmoCor. The difference between devices was also examined across age groups to investigate whether agreement changed over the PWV range associated with vascular aging.

The standard deviation of the paired differences was additionally interpreted according to the accuracy criteria recommended for PWV device-comparison studies [57].

6.9 Results

The validation results are presented at three levels: quality of the acquired pulse-wave signals, agreement between the PWV estimates obtained with the proposed system and SphygmoCor, and repeatability across repeated measurements. The agreement analysis was performed on the representative PWV value obtained for each participant and device according to the data-reduction procedure described in Section 6.8.2.

6.9.1 Waveform Quality and PWV Extraction

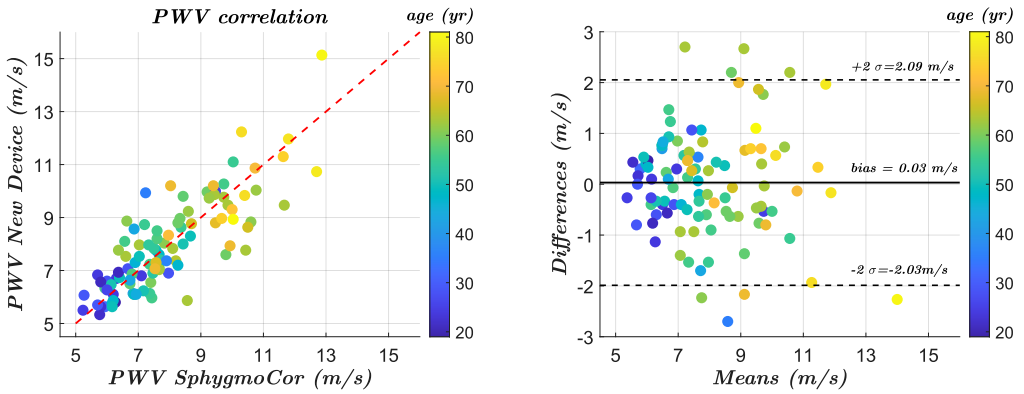
The wearable modules acquired periodic carotid and femoral pulse waves with sufficient stability and morphology for automatic fiducial-point extraction. The carotid waveform generally exhibited a more pronounced systolic upstroke, whereas the femoral waveform was typically smoother and more affected by the surrounding tissue and local sensor-coupling conditions. Despite these site-dependent differences, the processing procedure identified corresponding ITPs in the two waveforms and produced beat-level PTT estimates over the selected 10 s analysis windows.

The mean SNR was approximately 18.8 dB for the carotid recordings and 17.8 dB for the femoral recordings. These values complement the real-time visual assessment performed by the operators and indicate that the recorded waveforms contained a clearly detectable periodic pulse component.

The extracted PTT values were combined with the measured carotid–femoral distance according to Equation 6.2. The resulting PWV estimates covered a broad physiological range across the validation cohort, enabling comparison with the clinical reference system.

6.9.2 Agreement with the Reference System

Figure 6.8 summarizes the comparison between the proposed wearable system and SphygmoCor. The correlation plot in Figure 6.8a shows the PWV estimates obtained with the two devices for all 101 participants, with data points color-coded according to participant age. The measurements are distributed around the identity line, indicating that the proposed system reproduced the inter-subject PWV variation observed with the clinical reference.



(a) Correlation between PWV values obtained with the proposed system and SphygmoCor. Point color represents participant age.

(b) Bland–Altman comparison between the proposed system and SphygmoCor.

Fig. 6.8 Agreement between the proposed wearable system and SphygmoCor: (a) correlation plot of the PWV values obtained with the two devices; and (b) Bland–Altman plot showing the paired differences, mean bias, and limits of agreement. Adapted from [4].

The Pearson correlation coefficient was 0.823, while the Spearman rank-correlation coefficient was 0.809. Lin's concordance correlation coefficient was 0.823, and the intraclass correlation coefficient was 0.824. These results indicate a strong association and concordance between the two measurement systems.

The Bland–Altman analysis in Figure 6.8b showed a mean paired difference of $0.03\,\mathrm{m\,s^{-1}}$. The negligible bias indicates that the proposed system did not introduce a relevant systematic shift relative to SphygmoCor at cohort level. The standard deviation of the paired differences was $1.03\,\mathrm{m\,s^{-1}}$. Using the mean difference plus or minus two standard deviations, the lower and upper limits of agreement were $-2.03\,\mathrm{m\,s^{-1}}$ and $2.09\,\mathrm{m\,s^{-1}}$, respectively.

According to the accuracy criteria adopted for PWV device-comparison studies, the standard-deviation thresholds corresponding to good and acceptable accuracy at the measured bias were $0.694\,\mathrm{m\,s^{-1}}$ and $1.042\,\mathrm{m\,s^{-1}}$, respectively. The observed standard deviation of $1.03\,\mathrm{m\,s^{-1}}$ was therefore within the acceptable-accuracy range, although close to its upper limit [57].

Table 6.3 Age-stratified difference between the proposed wearable system and SphygmoCor.

Age group	n	**Mean difference** ($\mathrm{m\,s^{-1}}$)	**SD** ($\mathrm{m\,s^{-1}}$)
< 30 years	17	-0.19	0.56
30–49 years	14	0.01	1.04
50–69 years	58	0.09	1.11
≥ 70 years	12	0.03	1.03

6.9.3 Repeatability and Age-Stratified Performance

Measurement repeatability was evaluated from the triplicate acquisitions performed with each device. The mean within-subject standard deviation was $0.40\,\mathrm{m\,s^{-1}}$ for SphygmoCor and $0.63\,\mathrm{m\,s^{-1}}$ for the proposed wearable system. The prototype therefore exhibited greater within-subject variability than the clinical reference, while maintaining consistent PWV estimates across repeated acquisitions.

The difference between the proposed system and SphygmoCor was also evaluated across the four age groups defined in the validation protocol. The results are summarized in Table 6.3. The mean difference remained close to zero in every age group. The dispersion was lowest in participants younger than 30 years and higher in the three older groups.

The age-stratified results do not indicate a substantial systematic age-dependent bias. However, the subgroup estimates, particularly for the youngest and oldest participants, are based on relatively small samples and should therefore be interpreted cautiously.

6.10 Discussion

The validation results support the feasibility of the proposed wearable system for supervised one-step carotid–femoral PWV assessment [4]. Correlation and concordance coefficients of approximately 0.82 indicate that the system reproduced the inter-subject PWV variation measured by SphygmoCor. Moreover, the mean paired difference of $0.03\,\mathrm{m\,s^{-1}}$ shows that the complete measurement chain did not introduce a relevant systematic offset at cohort level.

The standard deviation of the paired differences was $1.03\,\mathrm{m\,s^{-1}}$, which remained within the acceptable-accuracy range adopted for PWV device-comparison studies, although close to its upper threshold. The negligible bias should therefore not be interpreted as complete equivalence between the two devices for every individual measurement. This is also reflected in the repeatability results: the mean within-subject standard deviation was $0.63\,\mathrm{m\,s^{-1}}$ for the proposed system and $0.40\,\mathrm{m\,s^{-1}}$ for SphygmoCor. The prototype thus showed greater acquisition-to-acquisition variability, plausibly related to repeated anatomical localization, fixation tension, sensor preload, and local tissue coupling.

The main contribution of the platform lies in the integration of the complete measurement workflow into two compact wearable modules. Compared with the rGO-R1 platform presented in the previous chapter, the commercial FSR provided a more robust and reproducible sensing element, while the digitally balanced bridge adapted the operating point to placement-dependent preload. Wearable fixation reduced continuous manual probe handling, and the hardware synchronization interface preserved the relative timing required for PTT estimation independently of BLE transmission delays. The resulting architecture enabled battery-powered, cable-free acquisition without sacrificing the timing information required for one-step PWV calculation.

The study has some limitations. The cohort excluded relevant cardiovascular conditions and was not balanced across age or PWV ranges; the results therefore cannot yet be generalized to pathological populations. Moreover, the proposed system and SphygmoCor could not be operated simultaneously or at identical locations, so short-term physiological variability may have contributed to the observed inter-device differences. The procedure also remained operator dependent: trained personnel were required to locate the arteries, adjust module position and fixation tension, repeat calibration, and verify waveform quality. Evaluation was limited to short, supervised recordings under predominantly static conditions. Motion artifacts, sensor displacement, packet loss, battery autonomy, antenna detuning, and end-to-end synchronization error were not characterized systematically. The hardware decoupling event provided a common timer start, but residual interrupt skew and oscillator drift remain unquantified and would become more relevant during prolonged recordings.

Future work should therefore address controlled-preload or multi-element sensing, automatic signal-quality assessment, motion robustness, long-term synchronization, wireless reliability, and evaluation in larger pathological cohorts before ambulatory use is considered.

6.11 Conclusions

This chapter presented a wearable and wireless system for one-step carotid–femoral PWV assessment. The platform integrates two FSR-based sensing modules, preload-adaptive analog conditioning, synchronized acquisition, BLE communication, rechargeable battery power, and host software for real-time waveform inspection and data storage. An offline processing workflow extracts carotid and femoral ITPs, estimates PTT, and calculates PWV from the measured propagation distance.

The system was evaluated against SphygmoCor Vx in a supervised study involving 101 volunteers. The comparison showed correlation and concordance coefficients of approximately 0.82, a negligible mean bias of $0.03\,\mathrm{m\,s^{-1}}$, and a standard deviation of the paired differences within the adopted acceptable-accuracy range. Repeatability was lower than that of the reference system but remained compatible with the observed cohort-level agreement.

The work therefore demonstrates the transition from the custom pressure-sensor proof of concept developed in the previous chapter to a complete PWV measurement platform quantitatively evaluated against a clinical reference. The proposed system reduces cabling and continuous manual probe handling while preserving the relative timing required for PTT estimation. Its current scope remains limited to short, supervised measurements performed by trained personnel; broader clinical validation, improved placement robustness, motion-tolerant acquisition, and long-term synchronization are required before ambulatory use can be considered.

Chapter 7

Conclusions

7.1 Main Outcomes of the Thesis

This thesis investigated the design of electronic acquisition systems starting from the information required by the application rather than from a predetermined sensing or sampling architecture. This perspective was applied primarily to multichannel *in vitro* neural recording and, in a complementary research line, to non-invasive arterial pulse-wave monitoring. Although the two domains involve different biological sources, interfaces, and observables, both required the complete measurement chain to be considered jointly, from signal coupling and analog conditioning to calibration, timing, embedded processing, communication, and experimental validation.

The principal contribution was the event-driven neural acquisition platform presented in Chapter 3. The work progressed from a modular low-noise analog front end and a motherboard for channel-specific threshold generation to two successive digital backends. The 15-channel MCU implementation established the complete acquisition workflow, whereas the 60-channel FPGA implementation enabled concurrent monitoring through parallel logic and a common time base. Under the adopted 1 ms time–channel representation, the MCU and FPGA implementations reduced the ideal continuous 12-bit payload at $10\,\mathrm{kS\,s^{-1}}$ by approximately 99.1% and 98.7%, respectively. The controlled validation further identified how SNR, threshold selection, spike amplitude, firing statistics, and calibration affected detection performance. The FPGA processing delay was approximately 80 ns, while the mean event-driven system latency ranged from approximately 0.34 ms to 1.45 ms and was

dominated by the analog threshold-detection path. The main outcome is therefore the demonstration that the time and spatial origin of neural activity can be retained with an approximately two-order-of-magnitude reduction in output payload when the experimental objective is event-centred network analysis. This advantage is obtained by intentionally discarding waveform morphology, and the platform is consequently an application-specific alternative rather than a universal replacement for continuous recording.

The cardiovascular studies extended the same system-level perspective to mechanical sensing. The rGO-R1 platform of Chapter 5 connected material development, flexible sensor geometry, mechanical coupling, bridge readout, embedded electronics, and physiological acquisition within one experimental system. It demonstrated pressure-dependent transduction, dynamic waveform acquisition, and preliminary carotid and femoral recordings. At the same time, it showed that material sensitivity alone is insufficient to establish suitability for physiological monitoring: deposition uniformity, preload, off-axis loading, mechanical recovery, electrical operating point, and robustness of the complete interface were equally decisive. The limitations identified in this platform informed the subsequent transition toward a more mechanically robust wearable architecture.

The wearable system of Chapter 6 progressed from waveform feasibility to the estimation of a clinically relevant parameter. Two FSR-based wireless modules acquired carotid and femoral pulse waves using a shared hardware timing reference established before placement. This separated synchronization from Bluetooth Low Energy transmission and preserved the relative timing required for Pulse Transit Time and Pulse Wave Velocity estimation. In a comparison with SphygmoCor involving 101 volunteers, the proposed system obtained correlation and concordance coefficients of approximately 0.82 and a mean paired difference of approximately $0.03\,\mathrm{m\,s^{-1}}$. The agreement supports the feasibility of simultaneous wearable carotid–femoral monitoring, although the lower within-subject repeatability relative to the reference system confirms that placement, coupling, and waveform-quality control remain central limitations.

7.2 Methodological Considerations

Across the thesis, the application-oriented design approach proved to be a strength. The relevant information was defined before selecting its electronic representation. For the neural platform, spike occurrence, channel identity, and relative timing were the target observables, allowing continuous waveforms to be replaced by a compact event representation. For PWV monitoring, continuous pulse-wave morphology had to be retained because the final measurement depended on fiducial-point extraction and on the relative delay between two sites. Data reduction was therefore not treated as an objective in itself; the objective was to preserve the minimum information required to answer the experimental question without imposing unnecessary acquisition, communication, or storage burdens.

A second general outcome is that the physical interface forms part of the measurement system. Electrode impedance, cell–electrode coupling, and channel-dependent noise influenced neural detectability, while tissue coupling, preload, anatomical location, and fixation influenced pulse-wave quality. The developed calibration procedures were consequently not secondary software functions, but mechanisms used to adapt the electronics to variable interfaces. Channel-specific neural thresholds, automatic bridge balancing, and site-dependent front-end adjustment all served this role. Their limitations also became visible: calibration can become unreliable when its observation interval is contaminated or when the operating conditions change after calibration.

Timing likewise emerged as an end-to-end property. Parallel FPGA logic provided a deterministic neural time base, but the complete latency was governed mainly by the preceding analog detector. In the wearable PWV system, wireless transfer was not used to establish synchronization; a direct hardware event aligned the local clocks before acquisition. These results show that temporal performance cannot be inferred from the sampling frequency or digital clock alone, but must be evaluated across the complete sensing, conditioning, synchronization, and processing path.

The validation strategy also reflected the maturity and practical constraints of each platform. Controlled electrical and synthetic tests provided known ground truth for the neural electronics; sensor-level and dynamic mechanical tests separated material and system behaviour in the rGO-R1 platform; and comparison with a clinical reference assessed the final PWV estimate of the wearable system. However,

some tests lacked established benchmark procedures, and in several cases rigorous validation was limited by the availability of biological preparations, calibrated reference instrumentation, repeated devices, time, and experimental resources. Ad hoc protocols were therefore necessary to answer specific engineering questions. This constitutes an important limitation: the evidence is not uniform across the thesis, and results obtained from preliminary or custom validation methods should not be interpreted as equivalent to standardized or clinical validation. At the same time, making these boundaries explicit allowed the experiments to remain aligned with the actual development stage and application objective of each system.

7.3 Limitations and Future Perspectives

The next development stage for the neural platform is a complete end-to-end 60-channel validation with neuronal cultures over repeated and extended recordings. This should include comparison with reference instrumentation, channel-dependent and non-stationary activity, calibration stability, packet integrity, power consumption, thermal behaviour, and connector reliability. Adaptive or quality-controlled threshold calibration, faster dedicated comparators, hybrid access to selected waveform segments, and integration of stimulation and local FPGA decision logic would extend the platform toward closed-loop experiments and larger arrays.

For pulse-wave monitoring, future work should improve the reproducibility of the mechanical interface. Controlled-force fixation, multi-element sensing, automated placement guidance, and real-time waveform-quality indices could reduce dependence on operator positioning and local anatomy. The wearable system should also be assessed in broader and more heterogeneous populations, including subjects with vascular disease, and through repeated or longitudinal protocols that evaluate day-to-day stability, motion robustness, battery operation, and unsupervised use.

More generally, future validation should move from application-specific laboratory protocols toward shared benchmarks and traceable reference procedures. Where no standard method exists, the test protocol, ground truth, uncertainty, and interpretation boundaries should be defined before data collection and reported as part of the instrument itself. This would strengthen comparability between prototypes and clarify which performance claims arise from controlled subsystem tests and which apply to the complete biological measurement.

The central conclusion of the thesis is that biomedical acquisition systems should be designed and evaluated as complete information paths rather than as collections of independent sensors, circuits, and algorithms. In the neural research line, this perspective enabled a compact event representation with approximately two orders of magnitude lower payload. In the cardiovascular line, it guided the progression from a custom pressure-sensitive material to a synchronized wearable system capable of estimating carotid–femoral PWV. The resulting platforms remain prototypes with application-specific limitations, but together they demonstrate how an information-centred and application-oriented methodology can support more efficient, accessible, and experimentally meaningful biological instrumentation.

References

[1] Multi Channel Systems MCS GmbH. *Standard 60MEA: Electrode Layout and Technical Specifications*. Multi Channel Systems MCS GmbH, Reutlingen, Germany, 2014. Includes the 60MEA200/30iR-Ti layout. Accessed: 2026-07-29.

[2] STMicroelectronics. *STM32H7 Nucleo-144 Boards (MB1364)*. STMicroelectronics, April 2026. User manual including NUCLEO-H723ZG. Accessed: 2026-07-29.

[3] Lattice Semiconductor. *MachXO2 Breakout Board Evaluation Kit User's Guide*. Lattice Semiconductor, February 2022. Accessed: 2026-07-29.

[4] Alessandro Sanginario, Marco Pogliano, Irene Buraioli, Marco Boscherini, Dario Leone, Alberto Milan, and Danilo Demarchi. Redefining vascular monitoring: A wearable with force sensor resistors for real-time pulse wave velocity assessment. *IEEE Transactions on Biomedical Engineering*, pages 1–10, 2026. Early access.

[5] Alessandro Sanginario, Irene Buraioli, Marco Boscherini, Stefania Vitale, Sabrina Conoci, Daniele Botto, Dario Leone, Alberto Milan, Artur Ciesielski, Paolo Samorì, and Danilo Demarchi. Reduced graphene oxide-based flexible pressure sensor for biomedical applications. *IEEE Sensors Journal*, 24(22):37090–37103, 2024.

[6] Marco Boscherini, Antonela Pjeci, Giulia Tomagra, Federico Picollo, and Alessandro Sanginario. A low-cost and ultra-low-data-rate event-driven platform for in vitro real-time multichannel neural signal acquisition. In *2026 IEEE 23rd Mediterranean Electrotechnical Conference (MELECON)*, pages 1–6. IEEE, 2026.

[7] Youngtae Nam and Bruce C. Wheeler. In vitro microelectrode array technology and applications. *Critical Reviews in Biomedical Engineering*, 39(1):45–61, 2011.

[8] Marie Engelene J. Obien, Kosmas Deligkaris, Torsten Bullmann, Douglas J. Bakkum, and Urs Frey. Revealing neuronal function through microelectrode array recordings. *Frontiers in Neuroscience*, 8:423, 2015.

[9] Andrew F. M. Johnstone, Guenter W. Gross, Dieter G. Weiss, Olaf H.-U. Schroeder, Anke Gramowski, and Timothy J. Shafer. Microelectrode arrays: A physiologically based neurotoxicity testing platform for the 21st century. *NeuroToxicology*, 31(4):331–350, 2010.

[10] Daniel A. Wagenaar, Radhika Madhavan, Jerome Pine, and Steve M. Potter. Controlling bursting in cortical cultures with closed-loop multi-electrode stimulation. *The Journal of Neuroscience*, 25(3):680–688, 2005.

[11] Alan L. Hodgkin and Andrew F. Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology*, 117(4):500–544, 1952.

[12] Carl Gold, Darrell A. Henze, Christof Koch, and György Buzsáki. On the origin of the extracellular action potential waveform: A modeling study. *Journal of Neurophysiology*, 95(5):3113–3128, 2006.

[13] Micha E. Spira and Aviad Hai. Multi-electrode array technologies for neuroscience and cardiology. *Nature Nanotechnology*, 8(2):83–94, 2013.

[14] Hernan G. Rey, Carlos Pedreira, and Rodrigo Quian Quiroga. Past, present and future of spike sorting techniques. *Brain Research Bulletin*, 119(Part B):106–117, 2015.

[15] Alberto Mazzoni, Federico D. Broccard, Elena Garcia-Perez, Pietro Bonifazi, Maria Elisa Ruaro, and Vincent Torre. On the dynamics of the spontaneous activity in neuronal networks. *PLOS ONE*, 2(5):e439, 2007.

[16] Jens-Oliver Muthmann, Hayder Amin, Evelyne Sernagor, Alessandro Maccione, Dagmara Panas, Luca Berdondini, Upinder S. Bhalla, and Matthias H. Hennig. Spike detection for large neural populations using high density multielectrode arrays. *Frontiers in Neuroinformatics*, 9:28, 2015.

[17] John D. Rolston, Robert E. Gross, and Steve M. Potter. A low-cost multielectrode system for data acquisition enabling real-time closed-loop processing with rapid recovery from stimulation artifacts. *Frontiers in Neuroengineering*, 2:12, 2009.

[18] Brett J. Kagan et al. In vitro neurons learn and exhibit sentience when embodied in a simulated game-world. *Neuron*, 110(23):3952–3969.e8, 2022.

[19] Md Sayed Tanveer, Dhruvik Patel, Hunter E. Schweiger, Kwaku Dad Abu-Bonsrah, Brad Watmuff, Azin Azadi, Sergey Pryshchep, Karthikeyan Narayanan, Christopher Puleo, Kannathal Natarajan, Mohammed A. Mostajo-Radji, Brett J. Kagan, and Ge Wang. Starting a synthetic biological intelligence lab from scratch. *Patterns*, 6(5):101232, 2025.

[20] Jan Müller, Marco Ballini, Paolo Livi, Yihui Chen, Milos Radivojevic, Amir Shadmani, Vijay Viswam, Ian L. Jones, Michele Fiscella, Roland Diggelmann, Alexander Stettler, Urs Frey, Douglas J. Bakkum, and Andreas Hierlemann.

High-resolution CMOS MEA platform to study neurons at subcellular, cellular, and network levels. *Lab on a Chip*, 15(13):2767–2780, 2015.

[21] Jelena Dragas et al. In vitro multi-functional microelectrode array featuring 59,760 electrodes, 2,048 electrophysiology channels, stimulation, impedance measurement, and neurotransmitter detection channels. *IEEE Journal of Solid-State Circuits*, 52(6):1576–1590, 2017.

[22] Reid R. Harrison and Cameron Charles. A low-power low-noise CMOS amplifier for neural recording applications. *IEEE Journal of Solid-State Circuits*, 38(6):958–965, 2003.

[23] Multi Channel Systems MCS GmbH. *MEA2100-System: Technical Specification*. Multi Channel Systems MCS GmbH, 2026. Undated online product documentation; year denotes access year. Accessed: 2026-07-28.

[24] Multi Channel Systems MCS GmbH. *MEA2100-Mini-System: Technical Specification*. Multi Channel Systems MCS GmbH, November 2024. Datasheet. Accessed: 2026-07-28.

[25] Axion BioSystems. *Maestro Multiwell Microelectrode Array Platform*. Axion BioSystems, 2026. Undated online product documentation; year denotes access year. Accessed: 2026-07-28.

[26] Michael Serra et al. A low-cost interface for multi-electrode array data acquisition systems. *BioTechniques*, 45(4):451, 453–454, 456, 2008.

[27] Yuan-Jen Chang, Lun-De Liao, Chin-Teng Lin, Hsin-Yi Lai, Ji-Lin Chen, Yu-Ting Yang, Ying-Chen Ting, Ya-Po Huang, Robby Wu, Nitish V. Thakor, and You-Yin Chen. A low-cost multi-electrode array system for the simultaneous acquisition of electrophysiological signal and cellular morphology. *Journal of Neuroscience and Neuroengineering*, 1(2):131–142, 2012.

[28] Luis D. Garma, Luca Matino, Guido Melle, Fabio Moia, Francesco De Angelis, Francesca Santoro, et al. Cost-effective and multifunctional acquisition system for in vitro electrophysiological investigations with multi-electrode arrays. *PLOS ONE*, 14(3):e0214017, 2019.

[29] Intan Technologies. *RHD2000 Evaluation System User Guide*. Intan Technologies, 2023. Product documentation, accessed 2026-07-28.

[30] Rodrigo Quian Quiroga, Zoltan Nadasdy, and Yoram Ben-Shaul. Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering. *Neural Computation*, 16(8):1661–1687, 2004.

[31] Michael Rizk and Patrick D. Wolf. Optimizing the automatic selection of spike detection thresholds using a multiple of the noise level. *Medical & Biological Engineering & Computing*, 47(9):955–966, 2009.

[32] Jan Müller, Douglas J. Bakkum, and Andreas Hierlemann. Sub-millisecond closed-loop feedback stimulation between arbitrary sets of individual neurons. *Frontiers in Neural Circuits*, 6:121, 2013.

[33] Maxwell Murphy, Stefano Buccelli, Yannick Bornat, David Bundy, Randolph Nudo, David Guggenmos, and Michela Chiappalone. Improving an open-source commercial system to reliably perform activity-dependent stimulation. *Journal of Neural Engineering*, 16(6):066022, 2019.

[34] Fabio Corradi and Giacomo Indiveri. A neuromorphic event-based neural recording system for smart brain-machine interfaces. *IEEE Transactions on Biomedical Circuits and Systems*, 9(5):699–709, 2015.

[35] Song Luan, Yan Liu, Ian Williams, and Timothy G. Constandinou. An event-driven SoC for neural recording. In *2016 IEEE Biomedical Circuits and Systems Conference (BioCAS)*, pages 404–407, 2016.

[36] Michele Cartiglia et al. A 4096 channel event-based multielectrode array with asynchronous output. *Nature Communications*, 15:7163, 2024.

[37] Joseph Negri, Vilas Menon, and Tracy L. Young-Pearse. Assessment of spontaneous neuronal activity in vitro using multi-well multi-electrode arrays: Implications for assay development. *eNeuro*, 7(1):ENEURO.0080–19.2019, 2020.

[38] STMicroelectronics. *STM32H723VE, STM32H723VG, STM32H723ZE, and STM32H723ZG: Arm Cortex-M7 32-Bit 550 MHz MCU, up to 1 MB Flash, 564 KB RAM, Ethernet, USB, 3x FDCAN, Graphics, and 2x 16-Bit ADCs*. STMicroelectronics, May 2025. Accessed: 2026-07-29.

[39] Lattice Semiconductor. *MachXO2 Family Data Sheet*. Lattice Semiconductor, March 2025. Accessed: 2026-07-29.

[40] Texas Instruments. *LMP7721 3-Femtoampere Input Bias Current Precision Amplifier*. Texas Instruments, December 2014. Datasheet, originally published January 2008 and revised December 2014. Accessed: 2026-07-29.

[41] Texas Instruments. *TLV854x 500-nA, RRIO, Nanopower Operational Amplifiers for Cost-Optimized Systems*. Texas Instruments, April 2018. Datasheet, originally published December 2016 and revised April 2018. Accessed: 2026-07-29.

[42] Analog Devices. *LT3042: 20 V, 200 mA, Ultralow Noise, Ultrahigh PSRR RF Linear Regulator*. Analog Devices, 2017. Data sheet. Accessed: 2026-07-29.

[43] Analog Devices. *LT3093: -20 V, 200 mA, Ultralow Noise, Ultrahigh PSRR Negative Linear Regulator*. Analog Devices, 2019. Data sheet. Accessed: 2026-07-29.

[44] Keysight Technologies. *B2900A Series Precision Source/Measure Unit*. Keysight Technologies, June 2020. Data sheet including the B2912A two-channel precision source/measure unit. Accessed: 2026-07-29.

[45] Giulia Tomagra, Federico Picollo, Alfio Battiato, Barbara Picconi, Silvia De Marchis, Alberto Pasquarelli, Paolo Olivero, Andrea Marcantoni, Paolo Calabresi, Emilio Carbone, and Valentina Carabelli. Quantal release of dopamine and action potential firing detected in midbrain neurons by multifunctional diamond-based microarrays. *Frontiers in Neuroscience*, 13:288, 2019.

[46] National Instruments. *PCI/PCIe/PXI/PXIe/USB-6259 Specifications*. National Instruments, August 2025. Specifications including the NI USB-6259. Accessed: 2026-07-29.

[47] Adam C. Snyder, Michael J. Morais, and Matthew A. Smith. Dynamics of excitatory and inhibitory networks are differentially altered by selective attention. *Journal of Neurophysiology*, 116(4):1807–1820, 2016.

[48] Ronald Hugh Barker. Group synchronizing of binary digital systems. In Willis Jackson, editor, *Communication Theory*, pages 273–287. Butterworths Scientific Publications, London, UK, 1953.

[49] Milton Friedman. The use of ranks to avoid the assumption of normality implicit in the analysis of variance. *Journal of the American Statistical Association*, 32(200):675–701, 1937.

[50] Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1):289–300, 1995.

[51] Maurice G. Kendall and B. Babington Smith. The problem of *m* rankings. *The Annals of Mathematical Statistics*, 10(3):275–287, 1939.

[52] Frank Wilcoxon. Individual comparisons by ranking methods. *Biometrics Bulletin*, 1(6):80–83, 1945.

[53] World Health Organization. *Global Report on Hypertension 2025: High Stakes: Turning Evidence into Action*. World Health Organization, Geneva, 2025.

[54] Stéphane Laurent, John Cockcroft, Luc Van Bortel, Pierre Boutouyrie, Cristina Giannattasio, Daniel Hayoz, Bruno Pannier, Charalambos Vlachopoulos, Ian Wilkinson, and Harry Struijker-Boudier. Expert consensus document on arterial stiffness: Methodological issues and clinical applications. *European Heart Journal*, 27(21):2588–2605, 2006.

[55] Yoav Ben-Shlomo, M. Spears, Clare Boustred, Margaret May, Simon G. Anderson, Emelia J. Benjamin, Pierre Boutouyrie, J. Kennedy Cruickshank, Stéphane Laurent, Carmel M. McEniery, et al. Aortic pulse wave velocity

improves cardiovascular event prediction: An individual participant meta-analysis of prospective observational data from 17635 subjects. *Journal of the American College of Cardiology*, 63(7):636–646, 2014.

[56] Raymond R. Townsend, Ian B. Wilkinson, Ernesto L. Schiffrin, Alberto P. Avolio, Julio A. Chirinos, John R. Cockcroft, Kevin S. Heffernan, Edward G. Lakatta, Carmel M. McEniery, Gary F. Mitchell, Samer S. Najjar, Wilmer W. Nichols, Elaine M. Urbina, and Thomas Weber. Recommendations for improving and standardizing vascular research on arterial stiffness: A scientific statement from the american heart association. *Hypertension*, 66(3):698–722, 2015.

[57] Bart Spronck, Dimitrios Terentes-Printzios, Alberto P. Avolio, Pierre Boutouyrie, Andrea Guala, Ana Jerončić, Stéphane Laurent, Eduardo C. D. Barbosa, Johannes Baulmann, Chen-Huan Chen, Julio A. Chirinos, Stella S. Daskalopoulou, Alun D. Hughes, Azra Mahmud, Christopher C. Mayer, Jeong Bae Park, Gary L. Pierce, Aletta E. Schutte, Elaine M. Urbina, Ian B. Wilkinson, Patrick Segers, James E. Sharman, Isabella Tan, Charalambos Vlachopoulos, Thomas Weber, Elisabetta Bianchini, and Rosa Maria Bruno. 2024 recommendations for validation of noninvasive arterial pulse wave velocity measurement devices. *Hypertension*, 81(1):183–192, January 2024.

[58] Keyu Meng, Xiao Xiao, Wenxin Wei, Guorui Chen, Ardo Nashalian, and Sophia Shen. Wearable pressure sensors for pulse wave monitoring. *Advanced Materials*, 34:2109357, 03 2022.

[59] Suvrajyoti Mishra, Smita Mohanty, and Ananthakumar Ramadoss. Functionality of flexible pressure sensors in cardiovascular health monitoring: A review. *ACS Sensors*, 7(9):2495–2520, 2022.

[60] Jonathan P. Mynard, Avinash Kondiboyina, Remi Kowalski, Michael M. H. Cheung, and Joseph J. Smolich. Measurement, analysis and interpretation of pressure/flow waves in blood vessels. *Frontiers in Physiology*, 11:1085, 2020.

[61] D. J. Korteweg. Ueber die fortpflanzungsgeschwindigkeit des schalles in elastischen röhren. *Annalen der Physik*, 241(12):525–542, 1878.

[62] A. Isebree Moens. *Die Pulscurve*. E. J. Brill, Leiden, 1878. 4 unnumbered pages, 145 pages, 1 unnumbered page; illustrations; 23 cm.

[63] Alberto P. Avolio, Mark Butlin, and Andrew Walsh. Arterial blood pressure measurement and pulse wave analysis—their role in enhancing cardiovascular assessment. *Physiological Measurement*, 31(1):R1–R47, 2010.

[64] Ji Jin, Xingguang Geng, Yitao Zhang, Haiying Zhang, and Tianchun Ye. Pulse wave analysis method of cardiovascular parameters extraction for health monitoring. *International Journal of Environmental Research and Public Health*, 20(3), 2023.

[65] Ramakrishna Mukkamala, Jin-Oh Hahn, Omer T. Inan, Lalit K. Mestha, Chang-Sei Kim, Hakan Töreyin, and Survi Kyal. Toward ubiquitous blood pressure monitoring via pulse transit time: Theory and practice. *IEEE Transactions on Biomedical Engineering*, 62(8):1879–1901, 2015.

[66] Guanqun Zhang, Mingwu Gao, Dan Xu, N. B. Olivier, and Ramakrishna Mukkamala. Pulse arrival time is not an adequate surrogate for pulse transit time as a marker of blood pressure. *Journal of Applied Physiology*, 111(6):1681–1686, 2011.

[67] Luc M. Van Bortel, Stéphane Laurent, Pierre Boutouyrie, Phil Chowienczyk, J. K. Cruickshank, Tine De Backer, Jan Filipovsky, Sofie Huybrechts, Francesco U. S. Mattace-Raso, Athanase D. Protogerou, Giuseppe Schillaci, Patrick Segers, Sebastian Vermeersch, and Thomas Weber. Expert consensus document on the measurement of aortic stiffness in daily practice using carotid-femoral pulse wave velocity. *Journal of Hypertension*, 30(3):445–448, 2012.

[68] Alberto Milan, Gaia Zocaro, Dario Leone, Francesco Tosello, Irene Buraioli, Domenica Schiavone, and Franco Veglio. Current assessment of pulse wave velocity: Comprehensive review of validation studies. *Journal of Hypertension*, 37(8):1547–1557, August 2019.

[69] Sofie A. M. Huybrechts, Daniel G. Devos, Sebastian J. Vermeersch, Dries Mahieu, Eric Achten, Tine L. M. de Backer, Patrick Segers, and Luc M. Van Bortel. Carotid to femoral pulse wave velocity: A comparison of real travelled aortic path lengths determined by mri and superficial measurements. *Journal of Hypertension*, 29(8):1577–1582, 2011.

[70] Koen Matthys and Pascal Verdonck. Development and modelling of arterial applanation tonometry: A review. *Technology and Health Care*, 10(1):65–76, 2002.

[71] John Allen. Photoplethysmography and its application in clinical physiological measurement. *Physiological Measurement*, 28(3):R1–R39, 2007.

[72] Francesco Scardulla, Gloria Cosoli, Susanna Spinsante, Angelica Poli, Grazia Iadarola, Riccardo Pernice, Alessandro Busacca, Salvatore Pasta, Lorenzo Scalise, and Leonardo D'Acquisto. Photoplethysmograhic sensors, potential and limitations: Is it time for regulation? a comprehensive review. *Measurement*, 218:113150, 2023.

[73] Benyu Jiang, Baoming Liu, Karen L. McNeill, and Philip J. Chowienczyk. Measurement of pulse wave velocity using pulse wave doppler ultrasound: Comparison with arterial tonometry. *Ultrasound in Medicine & Biology*, 34(3):509–512, 2008.

[74] J. M. Davies, M. A. Bailey, K. J. Griffin, and D. J. A. Scott. Pulse wave velocity and the non-invasive methods used to assess it: Complior, SphygmoCor, Arteriograph and Vicorder. *Vascular*, 20(6):342–349, 2012.

[75] Dilpreet Buxi, Jean-Michel Redouté, and Mehmet Rasit Yuce. Blood pressure estimation using pulse transit time from bioimpedance and continuous wave radar. *IEEE Transactions on Biomedical Engineering*, 64(4):917–927, 2017.

[76] Yang Yu, Andrew Lowe, Gautam Anand, Anubha Kalra, and Huiyang Zhang. The investigation of bio-impedance analysis at a wrist phantom with two pulsatile arteries. *Cardiovascular Engineering and Technology*, 14(6):810–826, 2023.

[77] Josep Solà, Andy Adler, Arnoldo Santos, Gerardo Tusman, Fernando Suárez Sipmann, and Stephan H. Bohm. Non-invasive monitoring of central blood pressure by electrical impedance tomography: first experimental evidence. *Medical & Biological Engineering & Computing*, 49(4):409–415, 2011.

[78] Martin Klein, Daniel Erni, and Dirk Rueter. Three-dimensional magnetic induction tomography: Practical implementation for imaging throughout the depth of a low conductive and voluminous body. *Sensors*, 21(22):7725, 2021.

[79] Jian Sun, Wei Zhuang, Gen Li, Gui Jin, Jia Xu, Ke Ma, Feng Wang, and Hua Feng. An experimental study of pulse wave measurements with magnetic induction phase shift method. *Technology and Health Care*, 26(S1):157–167, 2018.

[80] Takuya Sakamoto. Signal separation using a mathematical model of physiological signals for the measurement of heart pulse wave propagation with array radar. *IEEE Access*, 8:175921–175931, 2020.

[81] Chuyu Tang, Zhirong Liu, and Linlin Li. Mechanical sensors for cardiovascular monitoring: From battery-powered to self-powered. *Biosensors*, 12(8):651, 2022.

[82] Paolo Salvi, Giuseppe Lio, Carlos Labat, Enrico Ricci, Bruno Pannier, and Athanase Benetos. Validation of a new non-invasive portable tonometer for determining arterial pressure wave and pulse wave velocity: The PulsePen device. *Journal of Hypertension*, 22(12):2285–2293, 2004.

[83] R. Asmar, A. Benetos, J. Topouchian, P. Laurent, B. Pannier, A. M. Brisac, R. Target, and B. I. Levy. Assessment of arterial distensibility by automatic pulse wave velocity measurement: Validation and clinical application studies. *Hypertension*, 26(3):485–490, 1995.

[84] S. S. Hickson, M. Butlin, J. Broad, A. P. Avolio, I. B. Wilkinson, and C. M. McEniery. Validity and repeatability of the Vicorder apparatus: A comparison with the SphygmoCor device. *Hypertension Research*, 32:1079–1085, 2009.

[85] Zi-Bo Zhou, Tian-Rui Cui, Ding Li, Jin-Ming Jian, Zhu-Guo Li, Shou-Rui Ji, Xin Li, Jian-Dong Xu, Hou-Fang Liu, Yi Yang, and Tian-Ling Ren. Wearable continuous blood pressure monitoring devices based on pulse wave transit time and pulse arrival time: A review. *Materials*, 16(6):2133, 2023.

[86] Gabriel Chan, Rachel Cooper, Manish Hosanee, Kaylie Welykholowa, Panayiotis A. Kyriacou, Dingchang Zheng, John Allen, Derek Abbott, Nigel H. Lovell, Richard Fletcher, and Mohamed Elgendi. Multi-site photoplethysmography technology for blood pressure assessment: Challenges and recommendations. *Journal of Clinical Medicine*, 8(11):1827, 2019.

[87] Irene Buraioli, Davide Lena, Alessandro Sanginario, Dario Leone, Giulia Mingrone, Alberto Milan, and Danilo Demarchi. A new noninvasive system for clinical pulse wave velocity assessment: The Athos device. *IEEE Transactions on Biomedical Circuits and Systems*, 15(1):133–142, 2021.

[88] Matti Kaisti, Tuukka Panula, Joni Leppänen, Risto Punkkinen, Mojtaba Jafari Tadi, Tuija Vasankari, Samuli Jaakkola, Tuomas Kiviniemi, Juhani Airaksinen, Pekka Kostiainen, Ulf Meriheinä, Tero Koivisto, and Mikko Pänkäälä. Clinical assessment of a non-invasive wearable MEMS pressure sensor array for monitoring of arterial pulse waveform, heart rate and detection of atrial fibrillation. *npj Digital Medicine*, 2:39, 2019.

[89] Andrea Valerio, Irene Buraioli, Alessandro Sanginario, Dario Leone, Giulia Mingrone, Alberto Milan, and Danilo Demarchi. A new true wireless system for real-time pulse wave velocity assessment. *IEEE Sensors Journal*, 24(15):24365–24376, August 2024.

[90] Irene Buraioli, Stefania Vitale, Andrea Valerio, Alessandro Sanginario, Dario Leone, Sabrina Conoci, Artur Ciesielski, Alberto Milan, Danilo Demarchi, and Paolo Samorì. Graphene-based pressure sensor application in non-invasive pulse wave velocity continuous estimation. *Advanced Materials Technologies*, 10(5):2400876, 2025.

[91] Thu H. Huynh, Roozbeh Jafari, and Wan-Young Chung. A robust bioimpedance structure for smartwatch-based blood pressure monitoring. *Sensors*, 18(7):2095, 2018.

[92] Chenxi Yang, Yudi Dong, Yingying Chen, and Negar Tavassolian. A smartphone-only pulse transit time monitor based on cardio-mechanical and photoplethysmography modalities. *IEEE Transactions on Biomedical Circuits and Systems*, 13(6):1462–1470, 2019.

[93] T. Corina Margain, Emily Powell, Alexandra Clark, and Adam Bush. Phonocardiography based pulse wave velocity system for non-occlusive assessment of arterial stiffness. *Frontiers in Cardiovascular Medicine*, 12:1481836, 2025.

[94] C. B. Huang, S. Witomska, A. Aliprandi, M. A. Stoeckel, M. Bonini, A. Ciesielski, and P. Samori. Molecule–graphene hybrid materials with tunable mechanoresponse: highly sensitive pressure sensors for health monitoring. *Advanced materials*, 31(1):1804600, 2019.

[95] J. Hu, G. Dun, X. Geng, J. Chen, X. Wu, and T. L. Ren. Recent progress in flexible micro-pressure sensors for wearable health monitoring. *Nanoscale Advances*, 5(12):3131–3145, 2023.

[96] Chwee-Lin Choong, Mun-Bo Shim, Byoung-Sun Lee, Sanghun Jeon, Dong-Su Ko, Tae-Hyung Kang, Jihyun Bae, Sung Hoon Lee, Kyung-Eun Byun, Jungkyun Im, Yong Jin Jeong, Chan Eon Park, Jong-Jin Park, and U-In Chung. Highly stretchable resistive pressure sensors using a conductive elastomeric composite on a micropyramid array. *Advanced Materials*, 26(21):3451–3458, 2014.

[97] Hongju Wang, Lu Wang, Nannan Sun, Yang Yao, Liling Hao, Lisheng Xu, and Stephen E. Greenwald. Quantitative comparison of the performance of piezoresistive, piezoelectric, acceleration, and optical pulse wave sensors. *Frontiers in Physiology*, 10:1563, 2020.

[98] Ian B. Wilkinson, Carmel M. McEniery, Giuseppe Schillaci, Pierre Boutouyrie, Patrick Segers, Anne Donald, and Philip J. Chowienczyk. ARTERY society guidelines for validation of non-invasive haemodynamic measurement devices: Part 1, arterial pulse wave velocity. *Artery Research*, 4(2):34–40, 2010.

[99] Z. Wu, J. He, H. Yang, and S. Yang. Progress in aromatic polyimide films for electronic applications: Preparation, structure and properties. *Polymers*, 14(6):1269, 2022.

[100] Stefania Vitale, Hugo Puozzo, Shamil Saiev, Leïla Bonnaud, Antonio Gaetano Ricciardulli, Artur Ciesielski, David Beljonne, and Paolo Samorì. Tuning the piezoresistive behavior of graphene-polybenzoxazine nanocomposites: Toward high-performance materials for pressure sensing applications. *Chemistry of Materials*, 35(17):6909–6919, September 2023.

[101] Andrea De Marcellis, C. Reig, and M. D. Cubells-Beltrán. Current-based measurement technique for high-sensitivity detection of resistive bridges with external balancing through control voltages. *IEEE Sensors Journal*, 17(2):404–411, January 2017.

[102] Pyboard. Pyboard D-series hardware reference. Online hardware documentation, 2020. Last updated 11 May 2020. Available online: https://pybd.io/hw/pybd_sfxw.html;.

[103] Rosendo A. Rodriguez, Valerie Cronin, Timothy Ramsay, Deborah Zimmerman, Marcel Ruzicka, and Kevin D. Burns. Reproducibility of carotid–femoral pulse wave velocity in end-stage renal disease patients: Methodological considerations. *Canadian Journal of Kidney Health and Disease*, 3:109, 2016.

[104] Interlink Electronics. *FSR Integration Guide*. Interlink Electronics, 2026. Undated document, year denotes access year. Accessed: 2026-07-29.

[105] Cyrille Lebossé, Bernard Bayle, Michel de Mathelin, and Pierre Renaud. Nonlinear modeling of low cost force sensors. In *2008 IEEE International Conference on Robotics and Automation*, pages 3437–3442. IEEE, 2008.

[106] Cyrille Lebossé, Pierre Renaud, Bernard Bayle, and Michel de Mathelin. Modeling and evaluation of low-cost force sensors. *IEEE Transactions on Robotics*, 27(4):815–822, 2011.

[107] Interlink Electronics. *FSR 400 Series Data Sheet*. Interlink Electronics, 2026. Includes the FSR 402; undated document, year denotes access year. Accessed: 2026-07-29.

[108] STMicroelectronics. STM32WB55xx and STM32WB35xx: Multiprotocol wireless 32-bit MCU, Arm-based Cortex-M4 with FPU, Bluetooth 5.4 and 802.15.4 radio solution. Datasheet DS11929, Rev. 18, June 2026. Available online: https://www.st.com/resource/en/datasheet/stm32wb55cc.pdf.

[109] STMicroelectronics. *Guidelines for On-Board Antenna Reference Designs for STM32WB/STM32WBA MCUs*. STMicroelectronics, April 2026.

[110] Andrea Valerio, Irene Buraioli, Alessandro Sanginario, Giulia Mingrone, Dario Leone, Alberto Milan, and Danilo Demarchi. A region-based cross-correlation approach for tonometric carotid–femoral pulse wave velocity assessment. *Biomedical Signal Processing and Control*, 93:106161, 2024.

[111] Lisheng Xu, Shuran Zhou, Lu Wang, Yang Yao, Liling Hao, Lin Qi, Yudong Yao, Hongguang Han, Ramakrishna Mukkamala, and Stephen E. Greenwald. Improving the accuracy and robustness of carotid–femoral pulse wave velocity measurement using a simplified tube-load model. *Scientific Reports*, 12(1):5147, 2022.

[112] Lawrence I-Kuei Lin. A concordance correlation coefficient to evaluate reproducibility. *Biometrics*, 45(1):255–268, 1989.

[113] Terry K. Koo and Mae Y. Li. A guideline of selecting and reporting intraclass correlation coefficients for reliability research. *Journal of Chiropractic Medicine*, 15(2):155–163, 2016.

[114] J. Martin Bland and Douglas G. Altman. Statistical methods for assessing agreement between two methods of clinical measurement. *The Lancet*, 327(8476):307–310, 1986.

Appendix A

Supplementary Results

A.1 Neural Acquisition Platform

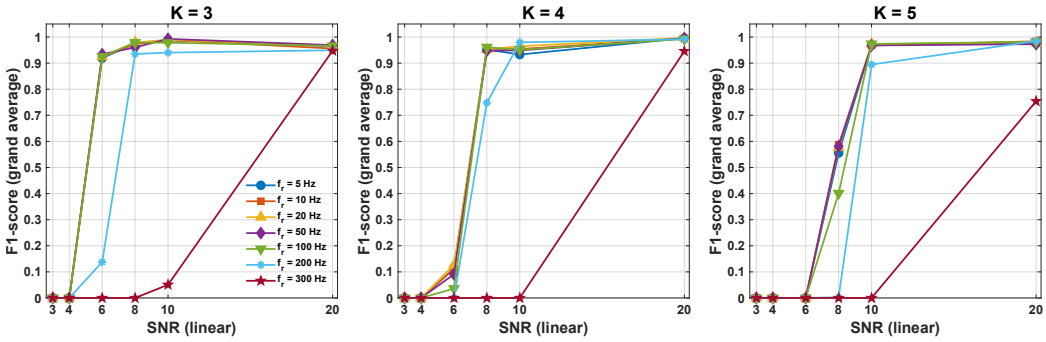
A.1.1 Complete FPGA Study 1 Results

This appendix reports supplementary results from FPGA Study 1. It includes alternative one-dimensional representations of the detection-performance dataset, complete condition-level detection metrics for each threshold multiplier, the conditions exceeding the repeatability criterion, the full statistical results for the effect of K, and the numerical temporal-jitter values. The experimental conditions and aggregation procedures are defined in Section 3.6.5.

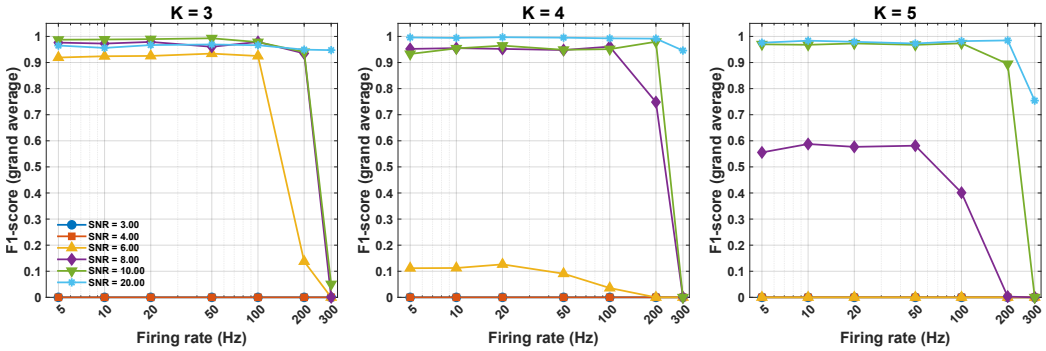
At $\mathrm{SNR}_{\mathrm{lin}} = 3$ and 4, F_1, precision, and recall were zero for all firing rates and threshold multipliers. These uniform-zero conditions are omitted from Tables A.1–A.3 to improve readability.

Alternative Detection-Performance Representations

Figure A.1 presents the same condition-level F_1 values reported by the operating-envelope heatmaps in Figure 3.10, plotted separately as functions of $\mathrm{SNR}_{\mathrm{lin}}$ and firing rate.



(a) Grand-average F_1 score as a function of $\mathrm{SNR}_{\mathrm{lin}}$, with firing rate as the curve parameter.



(b) Grand-average F_1 score as a function of firing rate, with $\mathrm{SNR}_{\mathrm{lin}}$ as the curve parameter.

Fig. A.1 Alternative representations of the FPGA Study 1 detection results. (a) Dependence of the F_1 score on $\mathrm{SNR}_{\mathrm{lin}}$, highlighting the progressive shift of the detection boundary as K increases. (b) Dependence on firing rate, showing the high-rate degradation under conditions in which the calibration interval was strongly contaminated by the regular spike train.

Complete Detection Metrics

Tables A.1–A.3 report the complete detection metrics for $K = 3$, $K = 4$, and $K = 5$, respectively. Values are the condition-level mean and standard deviation obtained using the aggregation procedure defined in the validation methodology.

Repeatability and Statistical Results

Table A.4 reports the conditions for which the inter-session coefficient of variation exceeded the predefined 10% threshold.

Table A.5 reports the SNR–firing-rate conditions in which the Friedman test identified a significant effect of K on the session-level F_1 score after false-discovery-rate correction.

Complete Temporal-Jitter Results

Table A.6 reports the temporal-jitter values used to generate Figure 3.12. The analysis was restricted to firing rates up to 100 Hz and to conditions with effective event detection.

Table A.1 Complete FPGA Study 1 detection results for $K = 3$.

SNR_{lin}	f_r (Hz)	F_1	Precision	Recall
6	5	0.919 ± 0.131	0.884 ± 0.017	0.963 ± 0.011
6	10	0.924 ± 0.127	0.893 ± 0.028	0.967 ± 0.007
6	20	0.925 ± 0.128	0.899 ± 0.012	0.962 ± 0.005
6	50	0.934 ± 0.118	0.919 ± 0.015	0.955 ± 0.015
6	100	0.925 ± 0.132	0.945 ± 0.009	0.912 ± 0.018
6	200	0.138 ± 0.202	0.772 ± 0.087	0.090 ± 0.027
6	300	0.000 ± 0.000	0.000 ± 0.000	0.000 ± 0.000
8	5	0.976 ± 0.129	0.978 ± 0.017	0.976 ± 0.014
8	10	0.973 ± 0.130	0.972 ± 0.010	0.973 ± 0.011
8	20	0.979 ± 0.122	0.987 ± 0.013	0.979 ± 0.011
8	50	0.960 ± 0.127	0.962 ± 0.013	0.960 ± 0.015
8	100	0.979 ± 0.102	0.979 ± 0.004	0.979 ± 0.004
8	200	0.935 ± 0.118	0.936 ± 0.037	0.935 ± 0.038
8	300	0.000 ± 0.000	0.000 ± 0.000	0.000 ± 0.000
10	5	0.987 ± 0.085	0.988 ± 0.011	0.987 ± 0.010
10	10	0.988 ± 0.103	0.988 ± 0.008	0.988 ± 0.007
10	20	0.990 ± 0.059	0.994 ± 0.014	0.989 ± 0.013
10	50	0.993 ± 0.071	0.998 ± 0.003	0.992 ± 0.007
10	100	0.978 ± 0.086	0.979 ± 0.014	0.978 ± 0.014
10	200	0.940 ± 0.215	0.955 ± 0.057	0.938 ± 0.067
10	300	0.051 ± 0.182	0.165 ± 0.155	0.041 ± 0.038
20	5	0.965 ± 0.074	0.997 ± 0.004	0.941 ± 0.034
20	10	0.956 ± 0.058	0.984 ± 0.020	0.933 ± 0.043
20	20	0.967 ± 0.056	0.997 ± 0.004	0.943 ± 0.041
20	50	0.969 ± 0.051	0.997 ± 0.003	0.946 ± 0.035
20	100	0.966 ± 0.049	0.997 ± 0.005	0.941 ± 0.029
20	200	0.949 ± 0.116	0.991 ± 0.006	0.919 ± 0.028
20	300	0.947 ± 0.150	0.948 ± 0.078	0.947 ± 0.079

Table A.2 Complete FPGA Study 1 detection results for $K = 4$.

SNR_{lin}	f_r (Hz)	F_1	Precision	Recall
6	5	0.112 ± 0.113	0.686 ± 0.233	0.063 ± 0.054
6	10	0.113 ± 0.099	0.891 ± 0.087	0.063 ± 0.041
6	20	0.126 ± 0.113	0.922 ± 0.068	0.072 ± 0.053
6	50	0.091 ± 0.084	0.927 ± 0.051	0.050 ± 0.027
6	100	0.036 ± 0.060	0.834 ± 0.091	0.019 ± 0.019
6	200	0.000 ± 0.000	0.000 ± 0.000	0.000 ± 0.000
6	300	0.000 ± 0.000	0.000 ± 0.000	0.000 ± 0.000
8	5	0.952 ± 0.128	0.932 ± 0.025	0.979 ± 0.001
8	10	0.955 ± 0.131	0.934 ± 0.017	0.979 ± 0.002
8	20	0.952 ± 0.132	0.933 ± 0.019	0.974 ± 0.006
8	50	0.948 ± 0.157	0.943 ± 0.013	0.963 ± 0.022
8	100	0.961 ± 0.132	0.957 ± 0.007	0.966 ± 0.010
8	200	0.748 ± 0.275	0.982 ± 0.002	0.662 ± 0.083
8	300	0.000 ± 0.000	0.000 ± 0.000	0.000 ± 0.000
10	5	0.933 ± 0.206	0.934 ± 0.059	0.934 ± 0.058
10	10	0.954 ± 0.135	0.957 ± 0.029	0.955 ± 0.025
10	20	0.965 ± 0.118	0.971 ± 0.030	0.965 ± 0.029
10	50	0.949 ± 0.163	0.954 ± 0.027	0.949 ± 0.029
10	100	0.951 ± 0.125	0.956 ± 0.018	0.951 ± 0.018
10	200	0.980 ± 0.103	0.979 ± 0.024	0.980 ± 0.024
10	300	0.000 ± 0.000	0.000 ± 0.000	0.000 ± 0.000
20	5	0.996 ± 0.028	0.997 ± 0.003	0.996 ± 0.004
20	10	0.994 ± 0.042	0.995 ± 0.004	0.994 ± 0.004
20	20	0.997 ± 0.011	0.998 ± 0.003	0.996 ± 0.004
20	50	0.995 ± 0.058	0.999 ± 0.001	0.995 ± 0.007
20	100	0.993 ± 0.059	0.995 ± 0.006	0.992 ± 0.008
20	200	0.992 ± 0.083	0.997 ± 0.006	0.990 ± 0.009
20	300	0.946 ± 0.165	0.952 ± 0.080	0.946 ± 0.086

Table A.3 Complete FPGA Study 1 detection results for $K = 5$.

$\mathrm{SNR}_{\mathrm{lin}}$	f_r (Hz)	F_1	Precision	Recall
6	5	0.000 ± 0.000	0.000 ± 0.000	0.000 ± 0.000
6	10	0.000 ± 0.000	0.000 ± 0.000	0.000 ± 0.000
6	20	0.000 ± 0.000	0.000 ± 0.000	0.000 ± 0.000
6	50	0.000 ± 0.000	0.000 ± 0.000	0.000 ± 0.000
6	100	0.000 ± 0.000	0.000 ± 0.000	0.000 ± 0.000
6	200	0.000 ± 0.000	0.000 ± 0.000	0.000 ± 0.000
6	300	0.000 ± 0.000	0.000 ± 0.000	0.000 ± 0.000
8	5	0.555 ± 0.219	0.971 ± 0.017	0.416 ± 0.127
8	10	0.588 ± 0.236	0.973 ± 0.021	0.455 ± 0.155
8	20	0.577 ± 0.225	0.967 ± 0.018	0.445 ± 0.135
8	50	0.581 ± 0.250	0.982 ± 0.002	0.451 ± 0.130
8	100	0.401 ± 0.264	0.979 ± 0.008	0.287 ± 0.112
8	200	0.003 ± 0.015	0.131 ± 0.129	0.001 ± 0.002
8	300	0.000 ± 0.000	0.000 ± 0.000	0.000 ± 0.000
10	5	0.970 ± 0.130	0.958 ± 0.009	0.982 ± 0.001
10	10	0.968 ± 0.130	0.955 ± 0.011	0.981 ± 0.002
10	20	0.973 ± 0.118	0.965 ± 0.009	0.982 ± 0.004
10	50	0.968 ± 0.132	0.960 ± 0.009	0.976 ± 0.005
10	100	0.974 ± 0.130	0.970 ± 0.005	0.977 ± 0.002
10	200	0.895 ± 0.201	0.982 ± 0.001	0.852 ± 0.058
10	300	0.000 ± 0.000	0.000 ± 0.000	0.000 ± 0.000
20	5	0.976 ± 0.130	0.976 ± 0.007	0.976 ± 0.007
20	10	0.984 ± 0.104	0.986 ± 0.005	0.983 ± 0.002
20	20	0.979 ± 0.112	0.982 ± 0.015	0.979 ± 0.013
20	50	0.973 ± 0.151	0.976 ± 0.009	0.973 ± 0.015
20	100	0.982 ± 0.130	0.989 ± 0.010	0.982 ± 0.022
20	200	0.985 ± 0.118	0.987 ± 0.014	0.984 ± 0.019
20	300	0.754 ± 0.373	0.926 ± 0.078	0.732 ± 0.141

Table A.4 FPGA Study 1 conditions exceeding the $CV = 10\%$ repeatability threshold.

K	$\mathrm{SNR}_{\mathrm{lin}}$	f_r (Hz)	CV (%)	$F_1^{(1)}$	$F_1^{(2)}$	$F_1^{(3)}$	$F_1^{(4)}$	$F_1^{(5)}$
3	6	200	29.6	0.186	0.162	0.084	0.111	0.145
3	10	300	94.2	0.100	0.000	0.000	0.086	0.067
4	6	5	79.1	0.253	0.067	0.144	0.049	0.045
4	6	10	61.1	0.201	0.098	0.166	0.047	0.050
4	6	20	66.7	0.260	0.072	0.160	0.067	0.073
4	6	50	51.4	0.153	0.053	0.128	0.069	0.051
4	6	100	95.7	0.095	0.027	0.028	0.020	0.008
5	8	5	21.7	0.736	0.592	0.555	0.436	0.457
5	8	10	24.6	0.797	0.535	0.671	0.490	0.445
5	8	20	20.4	0.758	0.581	0.599	0.489	0.457
5	8	50	20.1	0.782	0.550	0.537	0.476	0.561
5	8	100	31.2	0.609	0.388	0.378	0.359	0.272
5	8	200	154.2	0.003	0.010	0.000	0.001	0.000
5	20	300	17.2	0.947	0.645	0.827	0.673	0.678

Table A.5 FPGA Study 1 conditions showing a statistically significant effect of K on F_1. Reported p-values are false-discovery-rate corrected.

$\mathrm{SNR}_{\mathrm{lin}}$	f_r (Hz)	Preferred K	p	Kendall's W
6	5	3	0.00674	1.00
6	10	3	0.00674	1.00
6	20	3	0.00674	1.00
6	50	3	0.00674	1.00
6	100	3	0.00674	1.00
6	200	3	0.00674	1.00
8	5	3	0.01500	0.84
8	10	3	0.00674	1.00
8	20	3	0.01500	0.84
8	50	3	0.01500	0.84
8	100	3	0.00674	1.00
8	200	3	0.00674	1.00
10	5	3	0.02237	0.76
10	10	3	0.01500	0.84
10	50	3	0.01500	0.84
20	5	4	0.02237	0.76
20	10	4	0.00674	1.00
20	50	4	0.02237	0.76
20	200	4	0.02237	0.76
20	300	3	0.02237	0.76

Table A.6 Mean temporal jitter in FPGA Study 1 as a function of $\mathrm{SNR}_{\mathrm{lin}}$ and K. Values are mean ± standard deviation in milliseconds. A dash indicates that effective detection was not achieved.

$\mathrm{SNR}_{\mathrm{lin}}$	$K = 3$	$K = 4$	$K = 5$
6	0.239 ± 0.105	—	—
8	0.367 ± 0.104	0.252 ± 0.132	0.143 ± 0.134
10	0.347 ± 0.104	0.320 ± 0.129	0.199 ± 0.150
20	0.175 ± 0.130	0.168 ± 0.153	0.237 ± 0.163

Table A.7 Complete detection and temporal results for FPGA Study 2. Calibration was performed at $A_{\mathrm{ref}} = 150\,\mathrm{mV}$, and the resulting thresholds were held fixed throughout the amplitude sweep.

A_{spike} (mV)	$\mathrm{SNR}_{\mathrm{eff}}$	F_1	Precision	Recall	Jitter (ms)
50	3.33	0.570 ± 0.302	0.974 ± 0.119	0.461 ± 0.294	0.232 ± 0.116
75	5.00	0.938 ± 0.131	0.962 ± 0.118	0.921 ± 0.146	0.396 ± 0.083
100	6.67	0.950 ± 0.123	0.951 ± 0.123	0.950 ± 0.124	0.349 ± 0.151
125	8.33	0.958 ± 0.120	0.958 ± 0.120	0.959 ± 0.119	0.371 ± 0.095
150	10.00	0.969 ± 0.117	0.977 ± 0.069	0.969 ± 0.118	0.366 ± 0.091

A.1.2 Complete FPGA Studies 2 and 3 Results

This section reports the complete condition-level detection and temporal metrics for FPGA Studies 2 and 3. Values are presented as mean ± standard deviation according to the aggregation procedure defined in Section 3.6.7.

Study 2: Fixed-Calibration Amplitude Sweep

Table A.7 reports the detection performance obtained while reducing the spike amplitude after calibration. Calibration was performed once at $A_{\mathrm{ref}} = 150\,\mathrm{mV}$, with $\widehat{\sigma}_{\mathrm{cal}} = 15\,\mathrm{mV}$, and the resulting DAC thresholds were retained for all subsequent conditions. The effective linear SNR is therefore $A_{\mathrm{spike}}/\widehat{\sigma}_{\mathrm{cal}}$.

Study 3: Poisson-Distributed Spike Timing

Table A.8 reports the detection performance and temporal jitter obtained as a function of mean firing rate under Poisson-distributed spike timing.

Table A.8 Complete detection and temporal results for FPGA Study 3 under Poisson-distributed spike timing.

Mean f_r (Hz)	F_1	Precision	Recall	Jitter (ms)
5	0.935 ± 0.117	0.948 ± 0.084	0.930 ± 0.120	0.344 ± 0.024
10	0.933 ± 0.124	0.955 ± 0.085	0.922 ± 0.125	0.401 ± 0.026
20	0.927 ± 0.092	0.954 ± 0.062	0.907 ± 0.092	0.358 ± 0.037
50	0.903 ± 0.079	0.966 ± 0.025	0.853 ± 0.080	0.376 ± 0.034
100	0.855 ± 0.075	0.963 ± 0.034	0.773 ± 0.073	0.402 ± 0.024
200	0.779 ± 0.069	0.976 ± 0.023	0.652 ± 0.064	0.405 ± 0.033
300	0.639 ± 0.130	0.922 ± 0.091	0.497 ± 0.111	0.432 ± 0.039

A.1.3 Complete FPGA Study 4 Latency Results

Study 4 separated the latency of the registered FPGA processing path from that of the complete event-driven detection chain. The FPGA-internal latency is reported in nanoseconds, whereas the event-driven system latency is reported in microseconds.

For the event-driven system measurements, let μ_i and σ_i denote the event-level mean and standard deviation obtained within session i. The across-session quantities reported in the summary rows are

$$\overline{\mu}_{\text{sess}} = \frac{1}{N_{\text{s}}} \sum_{i=1}^{N_{\text{s}}} \mu_i, \tag{A.1}$$

$$s_\mu = \sqrt{\frac{1}{N_{\text{s}} - 1} \sum_{i=1}^{N_{\text{s}}} \left(\mu_i - \overline{\mu}_{\text{sess}}\right)^2}, \tag{A.2}$$

$$\sigma_{\text{RMS}} = \sqrt{\frac{1}{N_{\text{s}}} \sum_{i=1}^{N_{\text{s}}} \sigma_i^2}. \tag{A.3}$$

Here, $\overline{\mu}_{\text{sess}}$ is the mean of the session-level means, s_μ is their sample standard deviation, and σ_{RMS} summarizes the typical within-session event-level dispersion. The latter is distinct from the standard deviation across session means. For the FPGA-internal latency, $N_{\text{s}} = 15$; for each event-driven system-latency condition, $N_{\text{s}} = 10$.

FPGA-Internal Latency

Event-Driven System Latency

Inferential Analysis

The inferential analyses used the ten session-level mean latencies as paired observations. The SNR effect was evaluated separately for each value of K, while the threshold-multiplier effect was evaluated separately at each SNR.

The $\mathrm{SNR}_{\mathrm{lin}} = 20$ condition involved a comparator recovery transition rather than the intended leading transition. Its descriptive and inferential results are retained for completeness but should not be interpreted as physically equivalent to the measurements at $\mathrm{SNR}_{\mathrm{lin}} = 8$ and 10.

Table A.9 Per-session FPGA-internal latency results. The final row reports the mean and standard deviation of the 15 session means and the RMS of the 15 within-session standard deviations.

Session	**Mean** (ns)	**Within-session SD** (ns)	**Edges**
1	80.90	9.801	354
2	79.84	11.310	357
3	77.78	16.580	354
4	77.40	15.970	359
5	78.08	16.970	357
6	91.75	14.410	356
7	78.82	9.975	358
8	79.31	8.000	354
9	79.06	12.480	353
10	79.46	9.363	356
11	79.12	12.000	359
12	81.06	10.640	357
13	80.81	12.440	353
14	81.23	8.856	358
15	81.54	6.824	353
Across-session summary	$\mathbf{80.41 \pm 3.39}$	$\boldsymbol{\sigma_{\mathrm{RMS}} = 12.09}$	—

Table A.10 Per-session event-driven system latency for $K = 3$. The summary row for each SNR reports $\overline{\mu}_{\text{sess}} \pm s_{\mu}$ in the mean column and σ_{RMS} in the within-session SD column.

SNR_{lin}	Session	Mean (µs)	Within-session SD (µs)	Matched events
20	1	324	270	499
20	2	305	63	500
20	3	361	380	499
20	4	350	327	499
20	5	314	337	497
20	6	309	65	499
20	7	295	57	494
20	8	297	190	467
20	9	372	529	498
20	10	505	795	499
Across-session summary		$\mathbf{343.2 \pm 63.0}$	$\boldsymbol{\sigma_{\text{RMS}} = 373.8}$	—
10	1	740	465	496
10	2	615	193	492
10	3	716	364	499
10	4	641	159	499
10	5	710	594	495
10	6	694	178	488
10	7	686	133	499
10	8	716	551	499
10	9	717	441	497
10	10	779	506	499
Across-session summary		$\mathbf{701.4 \pm 46.7}$	$\boldsymbol{\sigma_{\text{RMS}} = 395.9}$	—
8	1	1143	326	495
8	2	927	163	499
8	3	1001	226	499
8	4	937	161	499
8	5	1007	366	499
8	6	941	176	483
8	7	1033	400	499
8	8	904	192	499
8	9	1194	606	499
8	10	965	145	499
Across-session summary		$\mathbf{1005.2 \pm 95.6}$	$\boldsymbol{\sigma_{\text{RMS}} = 309.8}$	—

Table A.11 Per-session event-driven system latency for $K = 4$. The summary row for each SNR reports $\overline{\mu}_{\text{sess}} \pm s_{\mu}$ in the mean column and σ_{RMS} in the within-session SD column.

SNR_{lin}	Session	Mean (µs)	Within-session SD (µs)	Matched events
20	1	269	68	500
20	2	334	464	478
20	3	354	406	498
20	4	339	318	496
20	5	309	79	500
20	6	375	86	499
20	7	385	341	498
20	8	304	59	499
20	9	367	298	498
20	10	435	601	499
Across-session summary		**347.1 ± 47.2**	**$\sigma_{\text{RMS}} = 327.0$**	—
10	1	1182	481	495
10	2	1044	405	498
10	3	1089	122	499
10	4	975	163	499
10	5	882	296	500
10	6	1092	359	499
10	7	1044	153	499
10	8	1008	103	499
10	9	1039	603	499
10	10	1071	540	495
Across-session summary		**1042.6 ± 79.2**	**$\sigma_{\text{RMS}} = 366.4$**	—
8	1	1458	84	440
8	2	1399	297	502
8	3	1428	101	496
8	4	1421	84	497
8	5	1421	86	484
8	6	1446	209	494
8	7	1442	88	483
8	8	1499	543	500
8	9	1430	87	484
8	10	1508	470	488
Across-session summary		**1445.2 ± 34.7**	**$\sigma_{\text{RMS}} = 263.6$**	—

Table A.12 Friedman tests evaluating the effect of SNR on the session-level mean event-driven system latency. The two tests, corresponding to $K = 3$ and $K = 4$, formed one predefined family ($m = 2$) for Benjamini–Hochberg false-discovery-rate correction.

K	χ^2_{Friedman}	p	Kendall's W	**Outcome**
3	20.000	< 0.0001	1.000	Significant
4	20.000	< 0.0001	1.000	Significant

Table A.13 Exact paired Wilcoxon signed-rank comparisons among SNR conditions using the session-level mean event-driven system latency. The six comparisons, comprising three SNR pairs for each value of K, formed one predefined family ($m = 6$) for Benjamini–Hochberg false-discovery-rate correction.

K	**Comparison**	p_{exact}	p_{FDR}	$\lvert r_{\text{rb}} \rvert$
3	SNR 20 vs. 10	0.002	0.002	1.000
3	SNR 20 vs. 8	0.002	0.002	1.000
3	SNR 10 vs. 8	0.002	0.002	1.000
4	SNR 20 vs. 10	0.002	0.002	1.000
4	SNR 20 vs. 8	0.002	0.002	1.000
4	SNR 10 vs. 8	0.002	0.002	1.000

Table A.14 Exact paired Wilcoxon signed-rank comparisons between $K = 3$ and $K = 4$ using the session-level mean event-driven system latency. The three comparisons, corresponding to the three SNR conditions, formed one predefined family ($m = 3$) for Benjamini–Hochberg false-discovery-rate correction.

SNR_{lin}	p_{exact}	p_{FDR}	$\lvert r_{\text{rb}} \rvert$	**Outcome**	**Direction**
20	1.000	1.000	0.000	n.s.	—
10	0.002	0.003	1.000	Significant	$K = 3 < K = 4$
8	0.002	0.003	1.000	Significant	$K = 3 < K = 4$