

Reliable Approximate Analog Computing in Flexible Electronics

Zur Erlangung des akademischen Grades einer

Doktorin der Ingenieurwissenschaften

von der KIT-Fakultät für Informatik des
Karlsruher Instituts für Technologie (KIT)

genehmigte
Dissertation

von

Paula Carolina Lozano Duarte

aus Bogotá, Kolumbien

Tag der mündlichen Prüfung: 28. Juli 2026

1. Referent: Prof. Dr. Mehdi B. Tahoori
Karlsruher Institut für Technologie (KIT)
2. Referentin: Prof. Dr. Jennifer Blain
Arizona State University (ASU)

Acknowledgements

Summarizing in just a few words what this doctoral journey has meant to me feels both difficult and deeply emotional. From an early age, I was lucky enough to have the unconditional support of my family and to grow up in an environment full of love and understanding, one that taught me I could achieve anything I set my mind to. I write these words thinking of my parents and my sister, but also of my grandparents and my aunt, who — without either of us fully realizing it — inspired me with her bravery to move to a country whose language I did not speak.

It is impossible not to think that none of this would have been possible without the support of my supervisor, Mehdi. I am deeply grateful for the trust he placed in me even before I joined the group, at a time when my background was perhaps not the most obvious fit for the thesis I would ultimately develop.

I would also like to acknowledge the public education system that has brought me to where I am today, and of which I am immensely proud to be a part. To the Universidad Pública de Navarra, for giving me the knowledge and skills that made it possible to complete this doctorate successfully. And to the Karlsruher Institut für Technologie, for providing the space and funding to carry out my research here, and for supporting my participation in conferences, workshops, and international research stays.

It is also beautiful to reflect on all the experiences that have shaped who I am and helped me grow: the scouts, who taught me the values of nonconformism and perseverance; theater, which taught me not to be afraid of making a fool of myself in public; and the student council, which showed me that even when the world seems to be against you, you will always find a group of people who, just like you, want to change the world and leave their mark. To all of you: thank you for being part of my life.

I could not forget my friends either, who from the very beginning have listened to my complaints and tears, and supported me unconditionally at every step of the way. This is also yours.

*“La vida no es la que uno vivió, sino la que uno recuerda
y cómo la recuerda para contarla.”*

— Gabriel García Márquez, *Vivir para contarla*

Abstract

Flexible Electronics (FE) based on Indium Gallium Zinc Oxide (IGZO) thin-film transistors offer a compelling platform for wearable healthcare, smart packaging, and large-area sensing: their substrates are mechanically compliant, their fabrication is low-temperature and low-cost, and their form factor adapts naturally to body surfaces and non-planar environments. Realizing the full potential of this technology, however, requires solving a set of interconnected design challenges that conventional silicon methodologies do not address. IGZO processes provide only n-type transistors, impose feature sizes of 600 nm or above, exhibit high device variability and defect densities, and rely exclusively on static power, which makes high-performance analog and mixed-signal circuit design fundamentally different from its silicon counterpart. This thesis presents a systematic investigation of analog and mixed-signal computing in FE, organized around five research directions that collectively span the design stack from individual circuit primitives to complete healthcare monitoring systems.

In the area of *analog function approximation*, a library of five Analog Building Blocks (ABBs) is developed and assembled into analog Kolmogorov-Arnold Networks (AKANs) via a hardware–software co-design framework that incorporates circuit-level error models into the training process. The resulting implementations achieve up to 125× area reduction over digital 8-bit splines and up to 55% area savings through hardware-aware pruning, with a maximum 5.5× improvement in the area×power product.

For *ADC co-design and co-optimization*, this thesis establishes that the ADC should be treated as a joint design variable with the downstream classifier rather than a fixed component. A novel Binary Search ADC architecture for FE achieves 2× area reduction over the state of the art, and in-training pruning and non-uniform quantization strategies reduce ADC hardware cost by up to 35.9× without compromising classification accuracy.

In *reliability and fault tolerance*, a hierarchical fault injection framework for IGZO Binary Search ADCs identifies that early conversion stages are disproportionately vulnerable, motivating a selective redundancy strategy that improves single-fault coverage from 60% to over 92% with only 4.2% area overhead. Two built-in self-test architectures—a distributed delay-based BIST achieving 93% defect coverage with 3% power overhead, and a spectral signature for fault detection occupying $6,390 \mu\text{m}^2$ at $16.7 \mu\text{W}$ —provide on-chip fault detection without external test equipment or analog-to-digital conversion of the test observable.

In *healthcare applications and clock generation*, the thesis integrates the preceding contributions into complete flexible systems for wearable and implantable healthcare monitoring. A cross-layer reliability framework combines a gain-boosted IGZO low-dropout regulator with integrated BIST, a fault-tolerant Flash ADC achieving up to 97.7% single-fault coverage, and a passive RF-based implantable sensor for hydrocephalus shunt monitoring that requires no on-board power. End-to-end co-design of flexible stress monitoring classifiers, evaluated across over 1200 configurations, identifies Pareto-optimal designs achieving 94% accuracy at $9\ \mu\text{W}$. The first analog feature extractors in FlexIC technology, combined with a NAS-inspired mixed-signal co-design framework, achieve $48\times$ area and $70\times$ energy-per-inference reductions over purely digital implementations.

Finally, the thesis presents the first phase-locked loop in n-type-only IGZO TFT technology, establishing bounded 1000 ppm clock stability at $115\ \mu\text{W}$ in a footprint of $0.009\ \text{mm}^2$ —more than $400\times$ lower power than prior FE clocking solutions—thereby addressing a system-level bottleneck that had been universally neglected in the FE literature.

Zusammenfassung

Flexible Elektronik (FE) auf Basis von Indium-Gallium-Zinkoxid (IGZO) Dünnschichttransistoren bietet eine vielversprechende Plattform für tragbare Gesundheitssysteme, intelligente Verpackungen und großflächige Sensorik: Die Substrate sind mechanisch nachgiebig, die Fertigung ist kostengünstig und bei niedrigen Temperaturen möglich, und die Formgebung passt sich natürlich an Körperoberflächen und nichtplanare Umgebungen an. Um das volle Potenzial dieser Technologie auszuschöpfen, müssen jedoch eine Reihe miteinander verknüpfter Entwurfsherausforderungen gelöst werden, die konventionelle Silizium-Methoden nicht adressieren. IGZO-Prozesse stellen ausschließlich n-Typ-Transistoren bereit, weisen Strukturgrößen von 600 nm oder mehr auf, zeigen hohe Bauteilvariabilität und Defektdichten und basieren ausschließlich auf statischer Leistungsaufnahme, was den analogen und gemischt-signal Schaltungsentwurf fundamental von seinem Silizium-Pendant unterscheidet.

Diese Dissertation präsentiert eine systematische Untersuchung analoger und gemischt-signalier Datenverarbeitung in der flexiblen Elektronik, gegliedert in fünf Forschungsrichtungen, die gemeinsam den gesamten Entwurfsstack von einzelnen Schaltungsprimitiven bis hin zu vollständigen Gesundheitsüberwachungssystemen abdecken.

Im Bereich der *analogen Funktionsapproximation* wird eine Bibliothek aus fünf analogen Grundbausteinen (ABBs) entwickelt und zu analogen Kolmogorov-Arnold-Netzwerken (AKANs) mittels eines Hardware-Software-Co-Design-Rahmens zusammengefügt, der schaltungsbasierte Fehlermodelle in den Trainingsprozess integriert. Die resultierenden Implementierungen erzielen eine bis zu 125× geringere Fläche gegenüber digitalen 8-Bit-Splines sowie bis zu 55 % Flächeneinsparung durch hardwarebewusstes Pruning, mit einer maximalen 5,5× Verbesserung im Fläche×Leistung-Produkt.

Für das *ADC-Co-Design und die Co-Optimierung* zeigt diese Dissertation, dass der ADC als gemeinsame Entwurfsvariable zusammen mit dem nachgelagerten Klassifikator betrachtet werden sollte und nicht als feste Komponente. Eine neue Binärsuch-ADC-Architektur für die FE erzielt eine 2× Flächenreduktion gegenüber dem Stand der Technik; Pruning während des Trainings sowie nicht-uniforme Quantisierungsstrategien reduzieren den ADC-Hardwareaufwand um bis zu 35,9×, ohne die Klassifikationsgenauigkeit zu beeinträchtigen.

Im Bereich *Zuverlässigkeit und Fehlertoleranz* identifiziert ein hierarchisches Fehlerinjektions - Framework für IGZO-Binärsuch-ADCs, dass frühe Konvertierungsstufen unverhältnismäßig stark gefährdet sind, was eine selektive Redundanzstrategie motiviert, die die

Einzelfehlerabdeckung von 60 % auf über 92 % bei nur 4,2 % Flächenmehraufwand verbessert. Zwei Built-In-Self-Test-Architekturen – ein verteilter verzögerungsbasierter BIST mit 93 % Defektabdeckung bei 3 % Leistungsmehraufwand und eine spektrale Signatur zur Fehlererkennung mit einem Flächenbedarf von $6\,390\,\mu\text{m}^2$ bei $16,7\,\mu\text{W}$ – ermöglichen die Fehlererkennung auf dem Chip ohne externes Testequipment oder Analog-Digital-Wandlung des Testmesswerts.

Im Bereich *Gesundheitsanwendungen und Taktgenerierung* integriert die Dissertation die vorangegangenen Beiträge in vollständige flexible Systeme zur tragbaren und implantierbaren Gesundheitsüberwachung. Ein schichtübergreifendes Zuverlässigkeits-Framework kombiniert einen verstärkungsgeboosteten IGZO-Niederausfallspannungsregler (LDO) mit integriertem BIST, einen fehlertoleranten Flash-ADC mit bis zu 97,7 % Einzelfehlerabdeckung und einen passiven RF-basierten implantierbaren Sensor zur Hydrozephalus-Shunt-Überwachung, der keine eigene Stromversorgung benötigt. Ein ganzheitliches Co-Design flexibler Stressüberwachungsklassifikatoren, ausgewertet über mehr als 1200 Konfigurationen, identifiziert Pareto-optimale Designs mit 94 % Genauigkeit bei $9\,\mu\text{W}$. Die ersten analogen Merkmalsextraktoren in der FlexIC-Technologie, kombiniert mit einem NAS-inspirierten gemischt-signalen Co-Design-Framework, erzielen $48\times$ geringere Fläche und $70\times$ geringere Energie pro Inferenz gegenüber rein digitalen Implementierungen.

Abschließend präsentiert die Dissertation die erste Phasenregelschleife (PLL) in ausschließlich n-Typ-IGZO-TFT-Technologie, die eine begrenzte Taktstabilität von 1000 ppm bei $115\,\mu\text{W}$ in einem Footprint von $0,009\,\text{mm}^2$ erreicht – mehr als $400\times$ geringere Leistungsaufnahme als bisherige FE-Taktlösungen – und damit einen systemseitigen Engpass beseitigt, der in der FE-Literatur bislang universell vernachlässigt wurde.

Zusammen demonstrieren diese Beiträge, dass zuverlässige, approximative und anwendungsfertige analoge und gemischt-signale Datenverarbeitung innerhalb der strengen Randbedingungen der flexiblen Elektronik realisierbar ist und dass das Co-Design über die Hardware-Software-Grenze hinweg der entscheidende methodische Ermöglicher ist.

Contents

Acknowledgements	i
Abstract	iii
Zusammenfassung	v
List of Figures	xiii
List of Tables	xxi
1. Introduction	1
1.1. Motivation	1
1.2. Research Contributions	2
1.3. Relation to Published Work	3
1.4. Thesis Organization	4
2. Background	5
2.1. Flexible Electronics	5
2.1.1. Indium Gallium Zinc Oxide Thin-Film Transistors	6
2.1.2. Design Constraints: Flexible Electronics vs. Silicon VLSI	7
2.1.3. Applications of Flexible Electronics	8
2.2. Analog-to-Digital Converters in Flexible Electronics	8
2.2.1. ADC Architectures and Trade-offs	9
2.2.2. Non-Uniform Quantization	10
2.3. Analog Function Approximation	10
2.3.1. Analog Building Blocks	11
2.3.2. Spline-Based Approximation	11
2.3.3. Kolmogorov-Arnold Networks	11
2.4. Approximate Computing	12
2.5. Reliability and Fault Tolerance in Flexible Electronics	13
2.5.1. Fault Models in IGZO Circuits	13
2.5.2. Fault Injection and Sensitivity Analysis	14
2.5.3. Built-In Self-Test for Mixed-Signal FE Circuits	14
2.5.4. Fault-Tolerant ADC Design Strategies	15
2.6. Clock Generation and Timing in Flexible AMS Systems	15
2.6.1. Voltage-Controlled Oscillators in IGZO	15
2.6.2. Phase-Locked Loops	15

2.7.	Machine Learning for Wearable Sensing	16
2.7.1.	Stress Monitoring as a Use Case	16
2.7.2.	Statistical Feature Extraction	16
2.7.3.	Classifier Architectures	17
2.7.4.	Hardware-Aware Training and Co-design	17
2.7.5.	Design Space Exploration	18
2.8.	Summary	18
I. Analog Function Approximation		21
Part Overview		23
3. Function Approximation Using Analog Building Blocks in Flexible Electronics		25
3.1.	Motivation and Problem Statement	25
3.2.	Analog Building Block Design	26
3.3.	Spline Implementation	27
3.3.1.	Bézier Spline Formulation	27
3.3.2.	Control Point Fitting	28
3.4.	KAN Assembly	29
3.5.	Experimental Evaluation	29
3.5.1.	Simulation Setup	29
3.5.2.	Spline Approximation Accuracy	30
3.5.3.	Comparison Against Digital Baseline	31
3.6.	Summary and Contributions	31
4. Co-Design of Kolmogorov-Arnold Networks for Analog Function Approximation		33
4.1.	Motivation	33
4.2.	Proposed Co-Design Methodology	34
4.2.1.	Hardware Architecture and Cost Characterization	34
4.2.2.	Error Modeling	35
4.2.3.	Pruning Strategy	35
4.2.4.	Hardware-Aware KAN Training	37
4.3.	Experimental Results	37
4.3.1.	Experimental Setup	37
4.3.2.	PVT Robustness Analysis	38
4.3.3.	Area and Power Results	39
4.3.4.	Comparison to Chapter 3 Baseline	42
4.4.	Summary and Contributions	42
II. ADC Co-design and Co-optimization		45
Part Overview		47

5. Design and In-Training Optimization of Binary Search ADC for Flexible Classifiers	49
5.1. Motivation	49
5.2. Proposed Binary Search ADC Architecture	50
5.2.1. Design constraints of the FE technology	50
5.2.2. Comparator design	50
5.2.3. Three-bit Binary Search ADC	51
5.3. In-Training Binary Search ADC Pruning	52
5.3.1. Motivation and key observation	52
5.3.2. Pruning methodology: NSGA-II-driven exploration	53
5.3.3. Quantization-aware training on pruned ADCs	54
5.3.4. Results	54
5.4. Summary and Contributions	56
6. Low-Power Binary Search ADC with High-Level Non-Uniform Input Optimization	57
6.1. Motivation	57
6.2. Non-Uniform Quantization Framework	58
6.2.1. Problem formulation	58
6.2.2. Artificial Bee Colony optimization	59
6.2.3. Threshold determination via K-Means	60
6.3. Four-Bit Non-Uniform Binary Search ADC Architecture	60
6.3.1. Circuit architecture	60
6.3.2. Non-uniform resistor ladder	61
6.3.3. Pruned non-uniform ADC	62
6.4. Post-Layout Validation Under PVT Conditions	63
6.4.1. Physical design	63
6.4.2. Trade-off scenarios	64
6.4.3. Comparison against Flash and SAR ADCs	65
6.4.4. PVT robustness	66
6.5. Results	67
6.5.1. Non-uniform ADC evaluation on biosignal datasets	67
6.5.2. Comprehensive comparison against the uniform baseline	69
6.5.3. Comparison with state-of-the-art non-uniform quantization	69
6.6. Additional Insights	70
6.7. Summary and Contributions	71
III. Reliability and Fault Tolerance	73
Part Overview	75
7. Fault Tolerant Design of IGZO-based Binary Search ADCs	77
7.1. Motivation	77

7.2.	Hierarchical Fault Injection Framework	77
7.2.1.	Overview	77
7.2.2.	Defect Modeling	78
7.2.3.	Fault Library and System-Level Classification	78
7.2.4.	Multi-Fault Scenarios	79
7.2.5.	Automated Simulation Environment	79
7.3.	Fault Sensitivity Analysis	80
7.4.	Vulnerability-Aware Fault-Tolerant Design	80
7.5.	Results	82
7.5.1.	System-Level Coverage	82
7.5.2.	Per-Component Analysis	82
7.5.3.	Multi-Fault Interaction Examples	83
7.6.	Summary and Contributions	83
8.	Distributed Delay-Based BIST for Mixed-Signal Circuits in Flexible Electronics	85
8.1.	Motivation	85
8.2.	Proposed RO-VCO Architecture	85
8.2.1.	Ring Oscillator	85
8.2.2.	Voltage-Controlled Oscillator	86
8.2.3.	Digital BIST Logic	86
8.3.	Experimental Setup	87
8.4.	PVT Characterization	88
8.5.	Defect Detection and BIST Output	89
8.5.1.	Comparison with the State of the Art	91
8.6.	Summary and Contributions	92
9.	Spectral Signatures for Parametric Fault Detection in Flexible Electronics	93
9.1.	Motivation	93
9.2.	Spectral Signature Methodology	94
9.2.1.	Signal Flow	94
9.2.2.	Ratiometric Normalization and Pass/Fail Decision	94
9.2.3.	Circuit Implementation	95
9.3.	Experimental Setup	96
9.4.	Results	96
9.4.1.	VCO Characteristics	96
9.4.2.	RMS Signatures Under Fault-Free PVT Conditions	97
9.4.3.	Normalized Signatures and Fault Coverage	98
9.4.4.	Sensitivity Analysis	100
9.4.5.	Comparison with ADC-Based Monitoring	100
9.5.	Summary and Contributions	101

IV. Healthcare Applications and Clock Generation	103
Part Overview	105
10. Cross-Layer Reliability for Flexible Healthcare Electronics	107
10.1. Motivation	107
10.2. Reliable Power Delivery: ReFlex-LDO	108
10.2.1. Architecture	108
10.2.2. DC and AC Performance	109
10.2.3. Structural Delay-Based BIST	111
10.2.4. Fault Coverage and Overhead	111
10.3. Fault-Tolerant Flash ADC	113
10.3.1. Architecture and Simulation Setup	113
10.3.2. Hierarchical Fault Injection Framework	114
10.3.3. Fault Sensitivity Analysis	115
10.3.4. Fault-Tolerant Comparator Variants	116
10.3.5. Single- and Multi-Fault Results	117
10.4. Passive Implantable Sensor	118
10.4.1. Clinical Motivation: Hydrocephalus Shunt Monitoring	118
10.4.2. RF-Based Sensing Principle	119
10.4.3. Prototype Fabrication and Experimental Validation	120
10.5. System Integration and Summary of Results	120
10.6. Summary and Contributions	121
11. Design Space Exploration of Low-Power Wearable Stress Monitoring Classifiers	123
11.1. Motivation	123
11.2. System Architecture	123
11.3. DSE Framework	124
11.4. Datasets	125
11.5. Results	126
11.6. Summary and Contributions	126
12. Co-Design for Mixed-Signal Smart Flexible Wearables for Healthcare	129
12.1. Motivation	129
12.2. Mixed-Signal System Architecture	129
12.3. Analog Feature Extractor Circuits	130
12.4. Flexible SAR ADC	131
12.5. Hardware-Aware Feature Selection via Differentiable Gating	132
12.6. Results	132
12.7. Summary and Contributions	133
13. Low-Power PLL-Based Clock Stabilization for Flexible IGZO AMS Systems	135
13.1. Motivation	135

13.2. Proposed PLL Architecture	135
13.2.1. Ring-Oscillator VCO	136
13.2.2. PFD and Charge Pump	136
13.3. System Integration and Target Specifications	137
13.4. Results	137
13.4.1. Comparison with Prior Clocking Approaches	140
13.5. Summary and Contributions	141
14. Conclusions and Future Work	143
14.1. Summary of Contributions	143
14.2. Open Problems and Future Research Directions	144
Bibliography	147
A. Acronyms and Symbols	165
B. List of Publications	167

List of Figures

2.1. PragmatIC's FlexIC design capabilities, integrating power management, analog/ sensor interfaces, processors, and communication modules into flexible ICs.	6
2.2. (a) Process schematic illustrating photolithographic patterning on flexible substrates, involving deposition of IGZO semiconductors and metallic interconnects on polyimide. (b) Cross-sectional structure of a unipolar IGZO TFT, illustrating the key layers.	7
2.3. Schematics of 3-bit ADC topologies: (a) Flash, (b) SAR, (c) $\Sigma\Delta$, (d) Binary Search, and (e) the resistor-NMOS comparator used in IGZO implementations.	9
2.4. Architectural comparison of a Multi-Layer Perceptron (MLP) and a Kolmogorov-Arnold Network (KAN). In the MLP, fixed activation functions reside at nodes; in the KAN, learnable spline activations are placed on edges.	12
3.1. Schematics of the five Analog Building Blocks: (a) inversion, (b) subtraction, (c) addition, (d) multiplication, and (e) squaring. All circuits use only N-type transistors and resistors compatible with the PragmatIC FlexIC IGZO process.	27
3.2. Architecture of the spline implementation from ABBs. The three additive sub-blocks (P_0 , $(P_1 - P_0)2x$, and $(P_0 - 2P_1 + P_2)x^2$) are summed to make a second-order Bézier spline.	28
3.3. Architecture of the analog KAN implementation. Multiple splines $B_0(x), \dots, B_i(x)$ are connected in parallel through resistors $R_0, \dots, R_i$ to realize a weighted sum at the common output node $\text{KAN}(x)$. A bias resistor R_{bias} sets the quiescent output level.	29
3.4. Layout view of a single spline.	30
3.5. Spline approximation accuracy for two representative target functions. Top row: ideal (dashed) vs. hardware-implemented (solid) spline output. The corresponding Normalized Mean Percentage Errors (NMPE) are included for both cases.	30
4.1. Block diagram of the proposed AKAN hardware–software co-design methodology. SPICE simulations characterize hardware splines and generate error models; KAN training incorporates these models alongside hardware costs; pruning reduces complexity at spline and network levels.	34

- 4.2. Spline architectures before and after pruning. (a) Original second-order Bézier spline. After setting individual coefficients to zero: (b) $k_0 = 0$, (c) $k_1 = 0$, (d) $k_2 = 0$. Each pruning step removes one or more ABB instances, yielding a direct reduction in area and power. 36
- 4.3. System-level robustness analysis across 27 PVT corners ($V_{DD} \in \{0.9, 1.0, 1.1\}$ V, Temp. $\in \{0, 27, 40\}^\circ\text{C}$, and FF/TT/SS process corners). The top row illustrates NMPE stability, demonstrating hardware-aware compensation, while the bottom row shows the total power consumption in log-scale. 39
- 4.4. Pareto distribution of NMPE versus AreaPower for the Iris Petal Length dataset approximating $\sin(x)$. The plot shows the baseline (3 nonpruned splines, star marker) and all possible pruning scenarios across the three splines in the network. Pruning configurations are labeled according to which coefficient (k_0 , k_1 , or k_2) is removed from each spline. Multiple configurations achieve lower NMPE and reduced Area $\times$ Power compared to the baseline, demonstrating favorable Pareto improvement through strategic parameter removal. 40
- 4.5. Transistor count breakdown of a printed MLP classification system for six UCI benchmark datasets. The ADC front-end (blue) dominates the total transistor budget in every case, confirming that ADC co-design is the primary lever for area and power reduction in FE classifiers [100]. 48
- 5.1. Proposed 3-bit Binary Search ADC [100]. (a) Block diagram with three cascaded comparison stages producing bits A2 (MSB), A1, and A0 (LSB). (b) Full circuit schematic: 5 comparators (COM0–COM4), a threshold amplifier (TA), and the 9-transistor control block. (c) Comparator schematic using N-type-only transistors, delivering both the comparison result and its complement. (d) Reference voltage waveforms at each comparison stage: constant midpoint at Stage 1, a two-level square wave at Stage 2, and a four-level staircase at Stage 3. 51
- 5.2. Pareto fronts of classification accuracy vs. normalized ADC area for six UCI benchmark datasets, comparing 2-, 3-, and 4-bit designs [100]. Pruned Binary Search ADCs (*Ours*) consistently achieve higher accuracy with lower area than the corresponding full-resolution designs (*Baseline*), confirming that partial quantization eliminates representations that are harmful rather than merely redundant. 55
- 6.1. Artificial Bee Colony (ABC) optimization framework for non-uniform ADC configuration. Employed, onlooker, and scout bees explore ADC level combinations. Each candidate is evaluated via quantization, MLP inference, and fitness computation. The best-performing configuration is selected and mapped to ADC hardware. 59

6.2.	Architecture of the non-uniform 4-bit Binary Search ADC [101]. The non-uniform resistor ladder (left) generates data-distribution-adapted reference voltages. The central schematic realizes the four-stage hierarchical comparison with 10 comparators, 8 inverters, and 26 pass gates. Reference voltage profiles per comparison stage are shown on the right, illustrating how the search window narrows non-uniformly across the input range.	61
6.3.	Comparison of 4-bit quantization schemes for a representative bimodal biosignal input [101]. Top: Uniform quantization divides the input range into 16 equal intervals regardless of data density, wasting resolution at the tails. Bottom: Non-uniform (data-driven) quantization places centroids near the two distribution peaks, preserving more classification-relevant information with the same number of levels.	62
6.4.	Example of a pruned 4-bit Binary Search ADC implementing non-uniform quantization for one biosensor channel [101]. Shaded components are absent from the physical implementation; only the branches corresponding to the retained quantization levels are instantiated, reducing both area and static power.	63
6.5.	Physical layout of the proposed Binary Search ADC designs [101]. (a) 3-bit design: the three comparison stages, control block, and non-uniform resistor ladder are visible as distinct layout regions. (b) 4-bit design: the additional fourth comparison stage and second control block are included, increasing total area by approximately $3\times$	63
6.6.	Area and power savings of the proposed Binary Search ADC relative to Flash and SAR ADC designs, for all trade-off scenarios at 3-bit (top) and 4-bit (bottom) resolution [101]. SC-A3 consistently achieves the best FoM, reducing area by up to $10.8\times$ and power by up to $3.8\times$ versus Flash.	66
6.7.	Accuracy deviation of the proposed Binary Search ADC under PVT conditions [101]. Deviation remains below 10% for all temperature and offset combinations up to 500 Hz under nominal supply, confirming robustness across the wearable biosignal frequency range. The $\pm 10\%$ V_{DD} variation (bottom panel) causes proportional power changes but does not significantly affect conversion accuracy.	67
6.8.	Pareto fronts illustrating the trade-off between classification accuracy and ADC levels for eight biosignal datasets [101]. Non-uniform ADC configurations (<i>This Work</i>) consistently achieve higher accuracy with fewer levels than the uniform 4-bit baseline, demonstrating the benefit of data-driven threshold placement across diverse biosignal modalities.	68
7.1.	Hierarchical fault injection workflow [96]. Left: five-step process from transistor-level characterization to performance evaluation with design hardening. Right: system-level propagation decision tree (fault is catastrophic if $DNL \geq 1$ LSB, marginal if $0.5 \leq DNL < 1$ LSB, benign otherwise) and the ADC schematic with selective redundancy elements highlighted.	78

7.2.	Automated fault simulation infrastructure [96]. A Python framework coordinates netlist parsing, resistive fault injection, parallel Spectre execution, and DNL-based metric extraction, enabling evaluation of more than 220 fault configurations per design variant.	79
7.3.	Fault sensitivity heat map of the baseline Binary Search ADC under single-fault injection [96]. Red: catastrophic ($DNL \geq 1$ LSB); yellow: marginal; green: benign. Fault vulnerability concentrates in early-stage elements (COM0, the control block's TA and T0-T1), while late-stage comparators COM3 and COM4 show 0% baseline coverage under multi-fault scenarios. . . .	80
7.4.	Physical layout comparison [96]. (a) Baseline COM3 comparator. (b) Fault-resilient COM3 with extended comparator-level redundancy (ECLR). (c) Complete baseline 3-bit Binary Search ADC. (d) Complete fault-resilient ECLR design, with 4.2% area overhead ($58,615 \mu m^2$ vs. $56,260 \mu m^2$).	81
8.1.	Schematics of (a) the proposed RO, (b) the VCO extending the RO with a single control transistor M0, (c) the digital BIST block using XNOR and flip-flop logic for stage-wise delay comparison, and (d) the fully integrated VCO+BIST with per-stage monitoring [95].	87
8.2.	Layout views of (a) the VCO, (b) the BIST digital logic, and (c) the integrated VCO+BIST design implemented in the PragmatIC FlexIC Gen3 process [95]. The BIST block occupies $4,585 \mu m^2$, adding 12% area relative to the VCO alone.	87
8.3.	Effect of V_{DD} on oscillation frequency and power per process-temperature corner [95]. Frequency decreases with increasing V_{DD} due to the resistive pull-up topology; the Fast corner achieves up to 122 kHz ($32.8 \mu W$) and the Slow corner is most power-efficient ($5 \mu W$).	89
8.4.	BIST output waveforms after full VCO+BIST integration [95]. Left: nominal operation (flag remains logic 0). Centre: OC defect detection at Stage 1 pull-up R5 (flags raised in Stages 1 and 3; Stage 2 flag analysis in Figure 8.5). Right: SC defect cases where the flag is not activated (inherently tolerant behavior explained below).	90
8.5.	Detailed VCO outputs for cases where the BIST fault flag is not activated [95]. The 2nd-stage SC fault is undetected because the asymmetric V_{DD} connection of the middle stage makes it inherently tolerant to this defect: the oscillation continues at the same frequency with a consistent inter-stage phase, producing no detectable phase deviation in the XNOR-based comparator.	91
9.1.	Complete signal flow of the proposed frequency-domain test signature circuit [97]. The CUT output bias drives the VCO; LPF and HPF extract spectral energy; normalized RMS signatures are compared against threshold τ to produce a binary pass/fail decision without any ADC or digital counter.	94
9.2.	Circuit-level implementation of the frequency-domain test signature circuit [97]. The IGZO-based VCO (M0-M4, R0-R10) drives first-order LPF and HPF filter networks. Component sizing: VCO R1-R7 = $172.47 \text{ k}\Omega$, $C = 0.3 \text{ pF}$; HPF $R = 251.27 \text{ k}\Omega$, $C = 2 \text{ pF}$; LPF $R = 172.47 \text{ k}\Omega$, $C = 2 \text{ pF}$	95

9.3.	Oscillation frequency of the proposed VCO as a function of V_{CONTROL} under fault-free conditions [97]. The monotonic characteristic across 250–350 kHz ensures an unambiguous mapping from CUT bias deviation to frequency shift, which is the foundational requirement of the spectral signature approach.	97
9.4.	RMS value of the second harmonic vs. V_{CONTROL} : nominal conditions (left) and all nine PVT corners (right) [97]. The absolute spread exceeds $5\times$ between extreme corners, but the ordering is preserved for all corners—the critical property enabling per-corner ratiometric normalization.	98
9.5.	Normalized RMS signatures S_{norm} under fault injection for all nine PVT corners at three operating points [97]. The green region indicates the detectable fault zone ($S_{\text{norm}} > \tau_{\text{high}}$) and the red region the undetectable zone. At $V_0 = 1.2$ V, slow-process corners clearly exceed τ_{high} for $ \Delta V \geq 150$ mV; at $V_0 = 2.0$ V, all corners remain within the pass band regardless of fault magnitude.	99
9.6.	Sensitivity of the normalized RMS signature to V_{CONTROL} deviations (left) and normalized RMS signatures across all PVT corners at $V_{DD} = 3$ V (right) [97]. Maximum sensitivity in the 1.2–1.5 V region yields 22.5% signature change for 300 mV faults vs. $\leq 9.5\%$ PVT spread, providing a $>12\%$ margin for reliable threshold placement.	100
9.7.	Physical layout of the proposed frequency-domain test signature circuit [97]. The VCO, HPF, and LPF blocks are visible; total area is $6,390 \mu\text{m}^2$. The compact footprint represents a $26\times$ area reduction versus the most compact ADC-based monitoring solution.	101
10.1.	Schematic and post-layout view of the proposed gain-boosted ReFlex-LDO [120]. The design includes a differential error amplifier (M5–M6), a second gain stage (M8–M9) that compensates for the limited intrinsic gain of IGZO TFTs, the pass transistor M_{pass} , a resistive feedback divider, and an on-chip output capacitor $C_0 = 0.95$ pF for dominant-pole compensation.	108
10.2.	Electrical performance of the ReFlex-LDO after post-layout extraction [120]. (a) V_{out} versus V_{DD} across PVT corners, showing line regulation. (b–c) V_{ref} versus temperature including 5000 Monte Carlo simulations. (d) Loop gain and phase response versus frequency after post-layout extraction. (e) PSRR versus frequency. (f) V_{out} response under supply ripple at V_{DD} . (g) Line transient response under varying V_{DD} conditions.	110
10.3.	3-bit Flash ADC and IGZO-TFT comparator used in the cross-layer reliability framework [120]. (a) Flash ADC architecture: seven comparators in parallel with a resistor ladder, followed by a 7-to-3 priority encoder. (b) IGZO-TFT comparator schematic: 7 transistors implementing the differential comparison and output latch.	113
10.4.	End-to-end hierarchical fault injection flow for IGZO-based Flash ADCs [120]. Single-fault injection and DNL classification build the vulnerability profile; multi-fault scenarios test design robustness; sensitive transistors guide targeted comparator hardening.	114

10.5. Proposed fault-tolerant comparator variants for the Flash ADC [120]. (a) D1: single redundant transistor M7 at the M3/M4 branch. (b) D2: single redundant transistor M8 at a complementary position. (c) D3: two redundant transistors (M7, M8) at the output latch (M5/M6). (d) D4: combined redundancy of D1, D2, and D3. (e) D5: a distinct configuration with two-level redundancy (M7–M12) providing the highest coverage for both OC and SC fault types.	116
10.6. Physical layout comparison of the baseline and D5 fault-tolerant Flash ADC [120]. (a) Baseline comparator (left) and full ADC (right): total area 0.484 mm ² . (b) D5 comparator (left) and full ADC (right): total area 0.500 mm ² (+3.1% area overhead, –8% power). The resistor ladder and comparators dominate the footprint; added transistors contribute marginally to total area.	118
10.7. Passive RF-based monitoring system for hydrocephalus shunt patency [120]. (a) System overview: external wearable transceiver interrogates the passively implanted antenna. (b) Valve closed (supine): sensing pads disconnected; antenna resonates at 915 MHz. (c) Valve open (upright): CSF electrically connects the pads; resonance shifts to 405 MHz. Persistence of 915 MHz resonance during upright posture indicates shunt malfunction.	119
11.1. Overview of the flexible real-time stress monitoring classification system. Biosensors capture physiological signals; per-sensor ADCs digitize them; feature extractors compute statistical features; the ML classifier predicts stress level.	124
11.2. Algorithmic flowchart of the proposed design space exploration. Feature selection, classifier training, pruning, and quantization compose a pipeline; hardware evaluation maps each design to its power consumption; Pareto-optimal accuracy–power solutions are extracted.	125
11.3. Resistor-NMOS layout of 2-input NAND gate optimized for 1 V operation, showing the pull-up resistor and n-type transistor arrangement.	125
11.4. Accuracy–power Pareto fronts from the full DSE across DT (a, d), SVM (b, e), and MLP (c, f) classifiers for the WESAD (a–c) and AffectiveROAD (d–f) datasets.	126
12.1. Overview of the proposed flexible mixed-signal system architecture. Analog feature extractors process all sensor signals concurrently in the analog domain; a MUX routes their outputs through a shared SAR ADC; digitized features are buffered for the digital MLP classifier.	130
12.2. (a) Analog feature extractor circuits: peak-detector for Max/Min, op-amp integrator for Mean, scaling amplifier for Sum. (b) Two-stage IGZO op-amp with differential pair, common-source amplifier, level-shifting, and buffer stage.	130
12.3. Schematic of the <i>n</i> -bit SAR ADC: sample-and-hold stage, <i>n</i> -bit SAR logic, R-2R ladder DAC, and gain-and-bias output stage.	131

12.4. Algorithmic overview of the hardware-aware co-design framework. Stochastic gating embedded in training jointly minimizes analog feature extraction cost and maximizes accuracy; post-training Lottery-Ticket pruning further reduces digital classifier area.	132
12.5. SPICE analog output (blue) versus ideal software-computed feature values (orange) for the SPD accelerometer x-axis: (a) Max, (b) Min, (c) Mean, (d) Sum. . .	133
13.1. Proposed PLL schematic. Left: PFD implemented from two pseudo-CMOS dynamic D flip-flops [38, 107]. Centre: charge pump and first-order RC low-pass filter. Right: three-stage ring-oscillator VCO with resistive pseudo-CMOS stages and V_{ctrl} -gated bias network.	136
13.2. Conceptual integration of the proposed PLL within a flexible AMS platform. A single PLL topology provides clock signals to multiple processing blocks at distinct frequency ranges; only the VCO load capacitor is adjusted for each target.	137
13.3. Open-loop VCO output frequency versus load capacitance C (log-log scale). Monotonic, near-linear relationship spanning 1 kHz to 300 kHz achieved by varying C from 0.27 pF to 52 pF.	138
13.4. Locked output frequency (top) and total power (bottom) across 27 PVT combinations: 3 temperatures $\times$ 3 process corners $\times$ 3 supply voltages. The PLL acquires lock reliably at 27 °C and 45 °C; at 0 °C IGZO mobility is insufficient for 300 kHz operation.	139
13.5. Absolute system power when each target is clocked by a SoA IoT CMOS PLL [91] (left bar) versus the proposed IGZO PLL (right bar). The system power (blue) is identical in both bars; only the clock contribution differs. Percentages above each bar indicate the clock's share of total power.	140

List of Tables

2.1. Design constraints in Flexible Electronics vs. Silicon VLSI	7
3.1. Hardware characteristics of the five Analog Building Blocks in the PragmatIC FlexIC process ($V_{dd} = 1.0\text{ V}$, $T = 27^\circ\text{C}$). Area and power are obtained from post-layout Cadence Spectre simulations.	27
3.2. Comparison of the analog KAN against digital spline implementations in the PragmatIC FlexIC process. Area and power are normalized to the 8-bit digital reference.	31
4.1. Summary of pruning impact on average and best NMPE across datasets, together with average area and power reduction.	41
4.2. Comparison of the pruned AKAN framework against the baseline ABB KAN from Chapter 3 on a common sigmoid benchmark. All metrics are normalized to the 8-bit digital spline reference.	42
5.1. Qualitative comparison of ADC architectures for FE (scale 1–4, where 1 is the lowest).	50
5.2. Full ADC comparison: area (μm^2) and power (nW) for 3- and 4-bit designs simulated at $V_{DD} = 1\text{ V}$, $V_{ref} = 1\text{ V}$, $T = 27^\circ\text{C}$, $f_{in} = 5\text{ Hz}$ [100].	52
5.3. Evaluation of printed MLP classification systems across ADC architectures and resolutions. Accuracy (%) is reported on the test set; area is in transistor count (TC) [100].	56
6.1. Simulation parameters used to evaluate robustness	64
6.2. Comparison of 3-bit and 4-bit design trade-off scenarios; FoM is the normalized product of area and power [101].	65
6.3. Comprehensive evaluation of non-uniform vs. uniform 4-bit ADC designs across biosignal datasets. B^1 denotes the uniform 4-bit ADC baseline [101].	69
6.4. Comparison of the non-uniform quantization approach of this work against the state-of-the-art pruning method [100]. Accuracy gain and level reduction are reported relative to the uniform 4-bit ADC baseline.	70
7.1. Fault coverage, area, and power for all design variants. Coverage: percentage of fault configurations with $\text{DNL} < 1\text{ LSB}$. Data from [96].	82
7.2. Multi-fault coverage by circuit component for all three design variants. Data from [96].	82

8.1.	Device properties for the RO, VCO, and BIST circuits. Data from [95].	88
8.2.	PVT corner analysis results for the RO. Data from [95].	89
8.3.	Comparison of proposed ROs and VCOs with FE state of the art. Data from [95].	92
9.1.	Fault coverage vs. fault magnitude $ \Delta V $ and operating point. Evaluated across 9 PVT corners $\times$ 2 fault polarities. Thresholds: $\tau_{\text{high}} = 1.10$, $\tau_{\text{low}} = 0.90$. Data from [97].	99
9.2.	Comparison with ADC-based flexible test circuits. Data from [97].	101
10.1.	Fault coverage of the ReFlex-LDO under structural delay-based BIST [120]. Open-circuit: 1 M Ω –1 G Ω ; short-circuit: 1–100 Ω	112
10.2.	Coverage, area, and power overhead versus number of BIST injection points [120].	112
10.3.	Baseline Flash ADC comparator and resistor ladder sizing [120].	114
10.4.	Fault sensitivity heat map for the 3-bit Flash ADC under single-fault injection. Values are RMS DNL deviations. Green: benign (<0.5); yellow: marginal (0.5–1.0); red: catastrophic (>1.0). Data from [120].	115
10.5.	Fault coverage, area overhead, and power change for all Flash ADC comparator variants. B = benign (DNL < 0.5); M = marginal (0.5–1.0); C = catastrophic (>1.0). Area and power values are relative to the baseline. Data from [120]. . .	117
13.1.	Target IGZO AMS systems for PLL integration ($V_{DD} = 3$ V).	137
13.2.	Comparison with prior clocking solutions in FE.	141

1. Introduction

1.1. Motivation

Flexible Electronics (FE) have emerged as a transformative platform for computing systems that must be lightweight, mechanically compliant, and fabricated at ultra-low cost [76, 20]. Built on polymer substrates such as polyimide using thin-film Indium Gallium Zinc Oxide (IGZO) transistors, FE devices can bend, stretch, and conform to irregular surfaces—capabilities that rigid silicon wafers cannot replicate [156]. This physical adaptability unlocks application domains inaccessible to conventional integrated circuits: disposable healthcare patches worn directly on the skin, smart packaging with embedded sensing and classification, large-area environmental monitoring systems, and conformable biomedical implants [68, 21, 115, 145, 139].

These applications share a common computational requirement: the ability to process analog sensor data, extract meaningful features, classify the result, and communicate or actuate—all within a severely constrained area and power budget, without external computation infrastructure [86, 166, 36, 112]. Realizing this processing pipeline on a single flexible substrate demands analog and mixed-signal (AMS) circuits that are compact, power-efficient, and robust to the elevated defect densities and device variability inherent in IGZO thin-film technology [96, 95, 120, 121].

Four fundamental challenges define the design space of AMS computing in FE. The *analog-to-digital interface* is the most expensive component in any FE classification system, often consuming 60–90% of the total transistor budget [6]; reducing this cost without degrading accuracy is the primary lever for system-level efficiency. The *reliability gap* between FE defect rates and the fault tolerance expected of a functioning sensor system requires targeted mitigation strategies absent from conventional CMOS design methodology [95, 96]. The *computation paradigm*—whether computation is best performed digitally after full digitization, or in the analog domain before it—remains an open question whose answer depends critically on the specific function to be evaluated, the available hardware resources, and the accuracy requirements of the application [99]. The *clock generation problem*—universally neglected in FE publications—is a fourth fundamental challenge: free-running ring oscillators suffer from unbounded PVT (Process-Voltage-Temperature) drift exceeding 30% frequency error in the kilohertz range, corrupting ADC sampling accuracy and forcing conservatively wide timing guard-bands [98].

This thesis addresses all four challenges across five research directions: analog function approximation, ADC co-design and co-optimization, reliability and fault tolerance, health-care wearable applications, and clock generation. Together, these contributions form a coherent body of work aimed at making fully integrated analog and mixed-signal computing on flexible substrates both technically feasible and practically deployable.

1.2. Research Contributions

The technical contributions of this thesis are organized into four parts, each corresponding to a distinct research thrust.

Part I: Analog Function Approximation. This part establishes a framework for performing nonlinear computation in the analog domain before digitization, reducing or eliminating ADC cost for functions that can be approximated with bounded error in continuous-time IGZO circuitry. Two contributions are presented: a library of five Analog Building Blocks (ABBs) and their composition into an analog KAN achieving $125\times$ area reduction over 8-bit digital splines; and a hardware–software co-design methodology (AKAN) that incorporates circuit-level error models into KAN training and applies hardware-aware pruning, achieving up to 55% area and 50% power reduction with an average $5.5\times$ improvement in the $\text{area}\times\text{power}$ product. A complementary published contribution—the holistic evolutionary framework for approximate Ternary Neural Networks (TNNs) in EGFET printed electronics [113]—co-optimizes the ADC interface and digital classifier, enabling battery-powered operation with a loss of accuracy of less than 5% this collaborative contribution is referenced for completeness; however, a detailed discussion is omitted since it lies beyond the central scope of this thesis.

Part II: ADC Co-design and Co-optimization. This part challenges the assumption that the ADC is a fixed component and demonstrates that joint optimization of the ADC and the downstream classifier yields substantial hardware savings. Two contributions are presented: a novel Binary Search ADC for FE achieving $2\times$ area reduction over the state of the art, paired with an NSGA-II-based in-training pruning methodology that reduces transistor count by $5\times$ on average; and a data-driven non-uniform quantization framework using an Artificial Bee Colony optimizer, achieving up to $12.5\times$ power and $35.9\times$ area reductions validated through post-layout simulation under PVT conditions. Two complementary published contributions extend this co-design principle to other classifier architectures: bespoke ADC–Decision Tree co-design [14] aligns Flash ADC comparison levels exactly with the classifier’s decision thresholds, eliminating all quantization levels unused by any decision boundary; and an in-training ADC pruning methodology for printed MLPs [6] reduces ADC transistor count by $5\times$ with less than 1% accuracy loss by identifying and removing redundant levels during the training loop. Both works are collaborative contributions to which the author contributed and are referenced throughout this thesis; nevertheless, their full methodological treatment lies beyond its primary focus.

Part III: Reliability and Fault Tolerance. This part addresses the elevated defect densities of IGZO manufacturing through targeted fault-tolerant design and built-in self-test strategies for analog and mixed-signal circuits. Three contributions are presented: a hierarchical fault injection and selective redundancy methodology for Binary Search ADCs that improves single-fault coverage from 60% to over 92% with an automated Python simulation pipeline; a distributed delay-based BIST achieving 93% defect coverage with 3% power overhead, exploiting a $1100\times$ area- and $5600\times$ power-reduced IGZO ring oscillator; and a frequency-domain ratiometric BIST that maintains fault separability across global PVT variation, occupying $6,390\ \mu\text{m}^2$ at $16.7\ \mu\text{W}$. A complementary published contribution addresses fault tolerance at the digital backend of Flash ADCs [97]: a lightweight sequential counter-based thermometer-to-binary encoder replaces area-intensive parallel bubble-error correction logic, achieving 91.4% average single-fault recovery across 4- to 7-bit resolutions with $2.7\times$ area and $3.9\times$ power savings over prior parallel approaches at 7-bit resolution; this collaborative contribution is included for contextual completeness, although its detailed analysis is outside the main scope of this dissertation.

Part IV: Healthcare Applications and Clock Generation. This part integrates the preceding contributions into complete flexible healthcare monitoring systems and addresses the overlooked system-level bottlenecks of reliable power delivery and clock generation. Four contributions are presented: a cross-layer reliability framework for flexible healthcare electronics combining a fault-tolerant Flash ADC with 97.7% single-fault coverage, a gain-boosted IGZO low-dropout regulator with integrated BIST, and a passive RF-based implantable sensor for hydrocephalus shunt monitoring; the first comprehensive design space exploration (DSE) of ML classifiers for stress monitoring in FE evaluating over 1200 designs and identifying Pareto-optimal flexible systems achieving 94% accuracy at $9\ \mu\text{W}$; the first analog feature extractors in FlexIC technology within a mixed-signal co-design framework that reduces area by $48\times$ and energy per inference by $70\times$; and the first PLL in n-type-only IGZO TFT technology providing bounded 1000 ppm clock stability at $115\ \mu\text{W}$.

1.3. Relation to Published Work

The contributions presented in this thesis are based on the following publications. Papers marked with an asterisk (*) are first-author contributions; all others are collaborative contributions to which the author made substantial technical contributions.

Part I (*) Function Approximation Using Analog Building Blocks in Flexible Electronics, ISQED 2025. (*) Co-Optimization of Analog Kolmogorov-Arnold Networks for Low-Power Function Approximation in Flexible Electronics, under review at IEEE Journal on Emerging and Selected Topics in Circuits and Systems (JETCAS), 2026.

Part II (*) Design and In-Training Optimization of Binary Search ADC for Flexible Classifiers, ASP-DAC 2025. (*) Low-Power Binary Search ADC with High-Level Non-

Uniform Input Optimization, under review at IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD), 2026.

Part III (*) Fault Tolerant Design of IGZO-based Binary Search ADCs, ISQED 2026. (*) Distributed Delay-Based BIST for Mixed-Signal Circuits in Flexible Electronics, ETS 2026. (*) Spectral Signatures for Parametric Fault Detection in Flexible Electronics, under review at ITC-Asia 2026.

Part IV Special Session: Reliable Emerging Electronics in Wearable and Implantable Healthcare Applications, IEEE VLSI Test Symposium (VTS), 2026. Exploration of Low-Power Flexible Stress Monitoring Classifiers for Conformal Wearables, ISLPED 2025. Feature-to-Classifer Co-Design for Mixed-Signal Smart Flexible Wearables for Healthcare at the Extreme Edge, ICCAD 2025. (*) Low-Power PLL-Based Clock Stabilization for Flexible IGZO AMS Systems, under review.

1.4. Thesis Organization

Chapter 2 establishes the technical background spanning FE technology, ADC architectures, analog function approximation, reliability and BIST, clock generation, and ML for wearable sensing. Parts I–IV (Chapters 3–13) present the research contributions in order of increasing system complexity. The thesis concludes with a summary of findings and directions for future research.

2. Background

This chapter establishes the technical foundations that underpin the contributions presented in this thesis. It begins with an overview of Flexible Electronics (FE) as a technology platform, covering materials, fabrication, and the design constraints that distinguish it from conventional silicon-based integrated circuits. It then introduces the analog and mixed-signal building blocks central to this work—analogue-to-digital converters, analog function approximation, and clock generation circuits—before discussing reliability challenges and testing strategies specific to FE. The chapter closes with an overview of the machine learning models targeted by the system-level co-design contributions.

2.1. Flexible Electronics

Flexible Electronics (FE) encompasses electronic devices designed to maintain functionality while bending, stretching, or conforming to non-rigid surfaces [76]. Unlike conventional integrated circuits built on rigid silicon wafers, FE systems are fabricated on mechanically compliant substrates such as polyimide, plastic films, or thin metallic foils [48, 17]. This mechanical adaptability opens a wide range of application domains inaccessible to rigid ICs, including wearable health monitors, conformable biomedical sensors, smart packaging, and large-area sensing systems [28].

The manufacturing process for FE adapts standard semiconductor fabrication techniques to these non-traditional substrates. Critical steps such as ion implantation and high-temperature annealing—routine in silicon processes—are eliminated or replaced with low-temperature alternatives [167]. Companies such as PragmatIC Semiconductor have developed platforms specifically optimized for FE: their FlexIC process defines feature sizes around 600 nm using photolithographic patterning of semiconductor and metallic interconnect layers directly on flexible substrates [125], as illustrated in Figure 2.1. The resulting process enables cost-effective, high-volume production compatible with standard IC supply voltages.




Figure 2.1: PragmatIC's FlexIC design capabilities, integrating power management, analog/sensor interfaces, processors, and communication modules into flexible ICs.

2.1.1. Indium Gallium Zinc Oxide Thin-Film Transistors

The active device in most FE platforms is the Indium Gallium Zinc Oxide (IGZO) Thin-Film Transistor (TFT) [156]. IGZO is an amorphous oxide semiconductor characterized by relatively high electron mobility, optical transparency, and compatibility with low-temperature deposition processes [122]. These properties make IGZO well-suited for large-area, low-power applications where the substrate cannot withstand the thermal budgets of conventional silicon processes [61].

Among the competing thin-film semiconductor materials—including hydrogenated amorphous silicon (a-Si:H), organic semiconductors, and low-temperature polycrystalline silicon (LTPS)—IGZO occupies a favorable middle ground [156]. Amorphous silicon suffers from low electron mobility (below $1 \text{ cm}^2/\text{V}\cdot\text{s}$) and pronounced instability under illumination and electrical stress [122]. Organic semiconductors offer excellent mechanical flexibility but are limited by even lower mobilities and high sensitivity to environmental conditions such as humidity and oxygen exposure [110]. LTPS achieves higher carrier mobilities but requires laser annealing steps that are both costly and difficult to scale to large substrate areas [28]. IGZO, by contrast, provides electron mobilities in the range of $10\text{--}50 \text{ cm}^2/\text{V}\cdot\text{s}$, good bias-stress stability, and process compatibility with substrate temperatures below 200°C , making it the preferred active layer for high-volume, large-area FE platforms [125].

A key limitation of IGZO technology is its inherently unipolar nature: only n-type devices are readily available [169]. The absence of complementary p-type transistors constrains circuit topologies, necessitating design strategies such as pseudo-CMOS, dynamic logic, or resistor-transistor logic (RTL) to implement functions that would otherwise rely on CMOS [156]. This restriction has direct implications for analog circuit design, where matched complementary pairs are routinely exploited for differential architectures, current mirrors, and amplifiers [63]. In IGZO circuits, such structures must be redesigned to function with n-type-only devices, typically at the cost of increased static power consumption [99].

IGZO TFTs also suffer from significant device variability. Process non-uniformities in carrier mobility, threshold voltage (V_{th}), and contact resistance—compounded by post-manufacturing changes arising from mechanical stress during bending—translate into large spreads in circuit performance [96]. These variability sources must be accounted for during both design and testing and motivate several of the fault-tolerance and reliability contributions presented in later chapters. Figure 2.2 shows a process cross-section and structural diagram of a typical unipolar IGZO TFT.




Figure 2.2.: (a) Process schematic illustrating photolithographic patterning on flexible substrates, involving deposition of IGZO semiconductors and metallic interconnects on polyimide. (b) Cross-sectional structure of a unipolar IGZO TFT, illustrating the key layers.

2.1.2. Design Constraints: Flexible Electronics vs. Silicon VLSI

The design space for FE circuits differs substantially from that of conventional silicon VLSI. Table 2.1 summarizes the key differences across the most relevant design dimensions [96].

Table 2.1.: Design constraints in Flexible Electronics vs. Silicon VLSI.

Aspect	Flexible Electronics	Silicon VLSI
Substrate	Plastic, polyimide	Rigid silicon wafers
Transistor type	IGZO (N-type only)	CMOS (N + P)
Feature size	≥ 600 nm	< 10 nm
Logic style	Pseudo-CMOS, unipolar	Full CMOS
Design focus	Area, power, simplicity	Power, performance, area
Analog precision	Limited by variability	High
Dominant power	Static ($>99\%$ at kHz rates)	Dynamic

FE is not intended to replace silicon-based platforms, but to complement them in domains where low cost, mechanical compliance, and large-area integration outweigh the need for high computational density [100]. The limited feature size and unipolar logic impose strict

upper bounds on integration density—typically a few thousand devices per chip—and restrict achievable analog performance. A particularly important consequence is that static power dominates total consumption (over 99% at sub-kHz operating frequencies), making transistor count a direct proxy for power consumption [145]. This motivates the pruning and co-design strategies throughout this thesis.

2.1.3. Applications of Flexible Electronics

The combination of mechanical flexibility, low-cost fabrication, and compatibility with non-planar surfaces positions FE as an enabling technology for several application domains:

- **Wearable health monitoring:** Conformable sensors and processing electronics placed directly on the skin enable continuous monitoring of physiological signals such as electrodermal activity, heart rate, and temperature [55, 120, 89].
- **Smart packaging and logistics:** FE enables disposable electronic tags with sensing, computation, and communication capabilities, supporting supply-chain monitoring and product authentication [104, 103, 7].
- **Environmental and industrial sensing:** Large-area FE arrays can cover surfaces to monitor structural health, gas concentrations, or soil parameters across spatial scales impractical for silicon ICs [56, 25, 29].
- **Biomedical implants:** Thin, flexible substrates reduce the mechanical mismatch between electronics and biological tissue, enabling conformable neural probes and cardiac patches [86, 22, 58].

Across these domains, a common requirement is the ability to perform on-sensor data processing to reduce communication bandwidth and energy expenditure. This motivates the integration of analog and mixed-signal computing blocks—including ADCs [5, 74], amplifiers [170, 111], feature extractors [139], and function approximators [94, 99]—directly on the flexible platform, which is the central design challenge this thesis addresses.

2.2. Analog-to-Digital Converters in Flexible Electronics

The analog-to-digital converter (ADC) is the critical interface between the physical, analog world and digital processing [84]. In FE, ADCs are indispensable for digitizing physiological signals, environmental measurements, or analog feature vectors before on-chip classification [101]. At the same time, they represent one of the largest sources of area

and power consumption in FE signal chains [14, 6], motivating their careful design and co-optimization with downstream digital blocks.

2.2.1. ADC Architectures and Trade-offs

Several ADC architectures have been explored in the FE context. Figure 2.3 illustrates the schematics of the principal topologies.




Figure 2.3.: Schematics of 3-bit ADC topologies: (a) Flash, (b) SAR, (c) $\Sigma\Delta$, (d) Binary Search, and (e) the resistor-NMOS comparator used in IGZO implementations.

Flash ADC. The Flash converter performs simultaneous comparisons of the input signal against $2^N - 1$ reference voltages generated by a resistor ladder, followed by a priority encoder [5]. This architecture offers the fastest conversion speed but requires an exponentially growing number of comparators with resolution, resulting in high power consumption and large area [137, 100]. Pruning techniques have been proposed to reduce comparator count with minimal accuracy loss [14, 6].

Successive Approximation Register (SAR) ADC. The SAR ADC employs a binary search algorithm: a digital-to-analog converter generates a reference corresponding to the current estimate, and a comparator refines the estimate bit by bit until convergence [45]. SAR ADCs achieve a good balance between speed, power, and resolution, but rely on precision capacitor arrays that are difficult to implement in IGZO due to the absence of complementary devices and limited component matching [8, 146].

Sigma-Delta ($\Sigma\Delta$) ADC. $\Sigma\Delta$ converters use oversampling and noise shaping with a 1-bit quantizer and digital filtering to achieve high resolution at low bandwidth [16]. While robust and accurate, the oversampling requirements and digital filter complexity make them poorly suited to the severely resource-constrained FE context [59].

Binary Search ADC. The Binary Search ADC represents a hybrid approach that combines aspects of Flash and SAR architectures [100]. It progressively halves the input voltage range at each conversion step, using a cascaded comparator tree that produces a binary output directly—without requiring a DAC feedback path or a digital encoder [96]. A 3-bit Binary Search ADC requires only 3 comparison stages and 5 comparators total, compared with 7 for an equivalent Flash converter. The sequential comparison structure, combined with a simple resistor ladder and an n-type-only comparator topology, makes this architecture particularly attractive for FE and is the primary ADC topology studied in this thesis [101].

2.2.2. Non-Uniform Quantization

Standard ADC architectures partition the input voltage range into 2^N equally spaced intervals, producing a uniform quantization grid. While simple to implement, uniform quantization is suboptimal when the probability distribution of the input signal is non-uniform—a condition common in physiological signals, where most measurements cluster in a restricted sub-range and large amplitudes are rare [72, 136]. Non-uniform quantization allocates quantization levels proportionally to the local density of the input distribution: more levels are placed in high-density regions, fewer in sparse ones. This principle, formalized independently by Max [108] and Lloyd [93], minimizes the mean squared quantization error for a given number of levels. The optimal threshold placement follows the Lloyd–Max conditions: each decision boundary lies at the midpoint between adjacent reconstruction levels, and each reconstruction level equals the conditional mean of the input within its quantization interval. In hardware, non-uniform quantization is realized by a non-uniform resistor ladder, where resistor values are sized proportionally to the desired interval widths. This modification requires no additional comparators and therefore incurs negligible area overhead relative to a conventional Flash or Binary Search ADC [100]. However, the optimal threshold placement depends on the input distribution and must therefore be determined jointly with the training of the downstream classifier—motivating the data-driven co-design methodology developed in Part II of this thesis.

2.3. Analog Function Approximation

Analog function approximation is the representation and evaluation of mathematical functions using continuous electrical signals, without analog-to-digital conversion and subsequent digital computation [94]. In FE, where converting sensor data to the digital domain

is costly in both area and power, analog preprocessing offers an attractive alternative: computations are performed in the analog domain before, or instead of, digitization [43]. Typical target functions include nonlinear activation functions for neural network inference, sensor calibration curves, signal compression, and statistical feature extraction [121, 139].

2.3.1. Analog Building Blocks

Analog Building Blocks (ABBs) are circuit primitives that implement fundamental mathematical operations—inversion, addition, subtraction, multiplication, and squaring—in the analog domain [31, 40]. By combining ABBs, more complex functions such as polynomial mappings, splines, and activation functions can be implemented in hardware [94]. In FE, each ABB must be designed to operate with n-type-only transistors and to be robust to device variability, making compactness and error resilience primary design criteria.

2.3.2. Spline-Based Approximation

Splines are piecewise polynomial functions that provide smooth approximations to arbitrary continuous functions. In this work, second-order Bézier splines are adopted as the approximation primitive [123]. A second-order Bézier spline is defined by three control points P_0 , P_1 , P_2 and evaluates to:

$$B(x) = P_0(1 - x)^2 + 2P_1(1 - x)x + P_2x^2, \quad (2.1)$$

which can be rewritten in a form amenable to analog implementation as:

$$B(x) = P_0 + (P_1 - P_0) 2x + (P_0 - 2P_1 + P_2) x^2. \quad (2.2)$$

This representation requires five ABB operations: two multiplications, one squaring, and two additions. The second-order spline strikes a practical balance between approximation accuracy and hardware cost; higher-order variants would improve fidelity but increase area and power proportionally [94].

2.3.3. Kolmogorov-Arnold Networks

Kolmogorov-Arnold Networks (KANs) are a class of neural network models grounded in the Kolmogorov-Arnold representation theorem, which states that any continuous multivariate function can be expressed as a finite superposition of continuous univariate functions [92, 70]. Unlike conventional multilayer perceptrons (MLPs), which fix the form of

activation functions and learn only linear weight matrices, KANs place learnable activation functions—implemented as splines—on network edges rather than nodes [49, 150]. A KAN layer computes:

$$\phi(x) = w_b b(x) + w_s \text{spline}(x), \quad (2.3)$$

where $b(x)$ is a basis function and $\text{spline}(x) = \sum_i c_i B_i(x)$ is a linear combination of Bézier basis functions with learnable coefficients c_i [94]. This formulation allows KANs to represent complex, smooth functions with substantially fewer parameters than equivalent MLPs, an advantage of particular importance in the resource-constrained FE setting. Figure 2.4 contrasts the architectural diagrams of a generic KAN and a standard MLP.




Figure 2.4.: Architectural comparison of a Multi-Layer Perceptron (MLP) and a Kolmogorov-Arnold Network (KAN). In the MLP, fixed activation functions reside at nodes; in the KAN, learnable spline activations are placed on edges.

The analog implementation of KANs in FE leverages the ABBs described above to realize each spline segment in hardware. The hardware–software co-design perspective—training the KAN subject to hardware constraints and then mapping it to physical circuits—is the central theme of the approximation contributions in Part I.

2.4. Approximate Computing

Approximate computing (AxC) is a design paradigm that deliberately accepts a bounded degree of computational error in exchange for improvements in area, power, or performance [13]. It is grounded in the observation that many application domains—including image processing, signal classification, and machine learning inference—are inherently error-tolerant: small deviations from exact results produce negligible degradation in application-level quality metrics such as classification accuracy or signal fidelity [114, 158].

Approximation can be introduced at multiple levels of the design hierarchy. At the **logic level**, approximate arithmetic units—adders, multipliers, comparators—are redesigned to

trade accuracy for reduced gate count or critical-path length [114, 138]. At the **data representation level**, reduced-precision fixed-point formats and non-uniform quantization schemes limit the number of bits used to represent intermediate values, directly reducing memory and arithmetic resources [93, 37]. At the **architecture level**, entire computational stages can be pruned or merged based on their sensitivity to the application output, guided by error propagation analysis [11, 159].

In the FE context, AxC takes on a particular importance. Because transistor count is directly proportional to static power consumption, any reduction in circuit complexity—even at the cost of bounded inaccuracy—yields proportional energy savings [162, 163, 165]. This tight area–power coupling means that the distinction between exact and approximate computing is not merely one of numerical fidelity, but one of system feasibility: an exact implementation may exceed the power budget of a flexible energy harvester, while an approximate one does not [145, 27]. The AxC principle therefore permeates several contributions of this thesis: non-uniform ADC quantization (Part II), hardware-aware pruning of KAN spline coefficients (Part I), and selective redundancy in fault-tolerant ADC design (Part III).

2.5. Reliability and Fault Tolerance in Flexible Electronics

Reliability is a first-order design concern in FE. The combination of high defect densities in thin-film processes, pronounced device-to-device variability, and ongoing mechanical stress during operation results in failure rates substantially higher than those observed in mature CMOS technologies [120, 95, 96].

2.5.1. Fault Models in IGZO Circuits

Faults in IGZO-based FE circuits arise from several distinct mechanisms:

- **Threshold-voltage shift:** Process non-uniformities cause V_{th} to deviate from its nominal value, altering comparator switching thresholds and ADC decision levels [145].
- **Open and short defects:** Incomplete deposition, pinhole formation, or metallic particle contamination during lithography can produce open-circuit or short-circuit faults in transistors and interconnects [96].
- **Mechanical stress-induced degradation:** Repeated bending cycles alter device parameters over time, leading to gradual drift of operating points and, in severe cases, cracking of interconnects [170, 122].

- **Contact resistance variation:** Poor ohmic contact formation between the IGZO channel and metallic source/drain electrodes introduces parasitic resistances that affect circuit gain and linearity [128, 122].

In mixed-signal circuits such as ADCs, these faults manifest as comparator misfires, incorrect digital output codes, or systematic offsets, all of which degrade conversion accuracy. A hierarchical view of fault propagation—from transistor level through comparator to full ADC conversion output—is essential for effective fault modelling and mitigation.

2.5.2. Fault Injection and Sensitivity Analysis

Fault sensitivity analysis identifies which components of a circuit contribute most to output errors under fault conditions. In the Binary Search ADC, errors in early conversion stages (close to the most significant bit) propagate through all subsequent decision levels and therefore have a disproportionately large impact on conversion accuracy [96]. Hierarchical fault simulation frameworks extend standard transistor-level injection with system-level propagation analysis, enabling efficient exploration of single and multi-fault scenarios across the conversion hierarchy [102]. Such frameworks, adapted to account for IGZO-specific failure modes, form the methodological basis of the fault-tolerance contributions in Part III.

2.5.3. Built-In Self-Test for Mixed-Signal FE Circuits

Built-In Self-Test (BIST) embeds test functionality on-chip, reducing or eliminating dependence on external Automatic Test Equipment (ATE). In FE, conventional ATE is impractical due to the absence of rigid packaging, the limited number of accessible pins, and the high device variability that makes parametric limits difficult to set globally [120].

BIST strategies explored in this thesis exploit two signal modalities. **Delay-based BIST** exploits ring oscillators (ROs) and voltage-controlled oscillators (VCOs) to generate time-domain signatures that are sensitive to device parameter variations. By embedding lightweight digital BIST logic within each oscillator stage, delay variations across stages can be monitored continuously to detect defects without external equipment [95]. **Frequency-domain solution** exploits the spectral response of the circuit under test (CUT) driven at its nominal input: radiometric RMS spectral signatures—which maintain fault separability across global PVT variations—provide robust pass/fail decisions under the large device variability of IGZO processes [97]. A critical advantage of both approaches over conventional ADC-based voltage monitoring is the avoidance of *test invalidation*: since the BIST circuitry is fabricated in the same variable process as the CUT, an ADC-based observer can itself produce incorrect pass/fail decisions under PVT variation, regardless of whether the CUT is actually faulty.

2.5.4. Fault-Tolerant ADC Design Strategies

Fault tolerance can be incorporated at the circuit level through selective redundancy and topology modification. Rather than applying uniform redundancy—which would double area and power—targeted strategies exploit the fault sensitivity analysis to identify critical components and protect only those [120].

2.6. Clock Generation and Timing in Flexible AMS Systems

Precise clocking is a prerequisite for synchronous AMS circuits, yet it is one of the most neglected design concerns in published IGZO systems. Free-running ring oscillators—the default timing source in FE—suffer from unbounded PVT drift: in the kilohertz operating range typical of FE AMS systems, moderate V_{th} shifts translate directly into frequency errors exceeding 30%, corrupting ADC sampling accuracy and forcing conservatively wide timing guard-bands [87, 81, 80, 127, 30]. Furthermore, when a state-of-the-art IoT CMOS PLL is integrated into representative IGZO systems, the clock alone can consume 29–89% of total system power [91], far exceeding the 25% clock-network share typical of silicon systems [53].

2.6.1. Voltage-Controlled Oscillators in IGZO

A Voltage-Controlled Oscillator (VCO) produces an output frequency that is tunable via control voltage. In FE, VCOs based on IGZO TFTs have been reported [80, 127, 87], but typically require high supply voltages (>20 V) and provide only coarse frequency tuning, limiting their integration into compact, low-voltage systems. On-chip VCOs that serve dual roles—as functional timing references and as distributed BIST monitor sensitive to circuit degradation—represent a more resource-efficient solution [46].

2.6.2. Phase-Locked Loops

A Phase-Locked Loop (PLL) provides reference-anchored frequency stability by closing a feedback loop around a VCO: a phase-frequency detector (PFD) compares the VCO output against a reference, a charge pump integrates the phase error, a loop filter smooths the control voltage, and the VCO tracks the filtered signal [91, 148, 129]. In silicon, PLLs are standard components offering parts-per-million frequency accuracy. In IGZO, the implementation of a PLL is challenged by the absence of p-type devices, limited carrier mobility,

and high V_{th} variability—constraints that preclude the direct reuse of silicon PLL topologies. The design of an adequate PLL for n-type-only IGZO, operating in the kilohertz range at 3 V with milliwatt-scale power, is one of the contributions of this thesis.

2.7. Machine Learning for Wearable Sensing

A key application driver for analog and mixed-signal computing in FE is the on-sensor execution of machine learning (ML) inference tasks [71, 68, 115]. Wearables and IoT edge nodes that monitor physiological signals must classify or extract features from raw sensor data, ideally without transmitting raw data to a cloud backend—saving energy and preserving privacy [75, 144].

2.7.1. Stress Monitoring as a Use Case

Psychological stress detection from wearable physiological signals is a representative and clinically relevant use case [1, 8, 139]. Electrodermal activity (EDA), also known as galvanic skin response (GSR), is a particularly suitable signal: it is recorded by passive skin-contact electrodes, operates in a low-frequency band (DC to a few hertz), and correlates with sympathetic nervous system arousal [164, 60]. The signal acquisition, feature extraction, and classification pipeline maps naturally onto the capabilities of FE platforms: low-frequency ADCs for digitization, analog or digital feature extractors, and compact classifiers [33, 150, 40, 4, 15]. This application context motivates the healthcare contributions of Part IV.

2.7.2. Statistical Feature Extraction

Raw physiological signals such as EDA operate at low frequencies (DC to a few hertz) and are sampled over windows of several seconds. Feeding raw time-domain samples directly to a classifier is impractical in FE: the sample count per window is large relative to the available transistor budget, and the temporal redundancy in slowly varying signals adds no discriminative value [1].

Feature extraction compresses each signal window into a compact vector of descriptive statistics that capture the essential characteristics of the signal while discarding redundant information. The statistical features most commonly used in wearable stress detection and well-suited to FE hardware are:

- **Mean (μ):** the average signal level over the window, characterizing the tonic component of EDA.

- **Maximum** and **Minimum**: the dynamic range of the signal, sensitive to large sympathetic responses.
- **Sum**: proportional to mean in fixed-length windows; useful as a proxy for total signal energy.
- **Variance** (σ^2): the second central moment, capturing phasic fluctuation amplitude.

These features are computable from simple arithmetic operations—addition, comparison, and squaring—that map directly onto the Analog Building Blocks described in Section 2.3.1 [94, 4]. Implementing feature extraction in the analog domain, before analog-to-digital conversion, therefore offers substantial savings in ADC resolution and transistor count. The design of the first analog feature extractors in FlexIC technology—and their integration into a mixed-signal co-design framework—is one of the contributions of Part IV.

2.7.3. Classifier Architectures

Several classifier architectures are considered in this thesis. **Multilayer Perceptrons (MLPs)** are fully connected neural networks with one or more hidden layers and fixed nonlinear activation functions. MLPs provide a good accuracy–complexity trade-off for tabular biosignal classification tasks and are the primary target for ADC co-design in Part II [101]. **Support Vector Machines (SVMs)** are margin-maximizing binary classifiers that map input features to a higher-dimensional space via a kernel function, well-suited to small-dataset regimes typical of physiological studies [1]. **Decision Trees (DTs)** are hierarchical binary classifiers that partition the feature space through threshold comparisons; their inference consists solely of comparisons, making them particularly amenable to mapping onto analog or mixed-signal hardware [14]. **Ternary Neural Networks (TNNs)** constrain weights to $\{-1, 0, +1\}$, reducing multiplication to addition and sign operations and enabling extremely compact hardware implementations in FE [113].

2.7.4. Hardware-Aware Training and Co-design

Standard ML training optimizes model accuracy without regard for the constraints of the target hardware platform. In the FE context, this leads to models that require ADC resolutions, analog precisions, or circuit areas that are infeasible [145, 14]. Hardware-aware training incorporates hardware cost models—including quantization noise, ADC resolution requirements, and circuit-level errors—directly into the training objective, enabling the joint optimization of model parameters and hardware specifications [11, 99, 163, 161]. A related paradigm that informs the feature selection methodology of this thesis is

Neural Architecture Search (NAS) [168, 44]. NAS automates the design of neural network architectures by treating structural decisions—layer count, width, connectivity—as learnable variables within an outer optimization loop. Early NAS methods relied on reinforcement learning or evolutionary search over discrete architecture spaces, incurring prohibitive computational cost [168]. Differentiable NAS (DNAS) addresses this by relaxing discrete choices into continuous probability distributions over candidate operations, enabling gradient-based joint optimization of architecture and weights in a single training pass [44]. In this thesis, a NAS-inspired differentiable gating mechanism is embedded directly into MLP training to jointly select which analog feature extractors to activate and how to weight their outputs, thereby minimizing analog hardware cost without a separate architecture search phase. This co-design principle—that hardware and software layers of a system must be optimized together to unlock the full efficiency potential of FE platforms—runs throughout the ADC co-design contributions of Part II and the application contributions of Part IV.

2.7.5. Design Space Exploration

Design Space Exploration (DSE) systematically surveys the trade-off frontier between competing design objectives. In a wearable stress monitoring system, the relevant axes include classifier accuracy, ADC resolution, feature vector length, total power, and silicon area. A DSE methodology maps feasible design points in this space and identifies Pareto-optimal configurations—those for which no improvement in one objective can be achieved without degrading another [1, 139, 101, 99].

2.8. Summary

This chapter introduces the five technical pillars that underpin the thesis contributions:

1. **Flexible Electronics and IGZO TFTs:** A mechanically compliant, n-type-only thin-film technology with unique design constraints arising from unipolar logic, large device variability, static-power dominance, and feature sizes around 600 nm.
2. **ADC architectures for FE:** Flash, SAR, $\Sigma\Delta$, and Binary Search topologies, with the Binary Search ADC emerging as the preferred balance of area, power, and hardware simplicity for low-frequency FE applications.
3. **Analog function approximation:** Analog Building Blocks, Bézier splines, and Kolmogorov-Arnold Networks as a framework for performing nonlinear computation in the analog domain, avoiding costly digitisation.

4. **Reliability and BIST:** Fault models specific to IGZO, hierarchical fault simulation, selective redundancy strategies, and delay- and frequency-domain BIST techniques for in-field self-testing of FE mixed-signal circuits without the risk of test invalidation.
5. **ML for wearable sensing:** MLP, SVM, DT, and TNN classifiers for on-sensor physiological data processing, and the hardware-aware co-design methodology that connects model training to physical circuit constraints.

Each subsequent part of this thesis builds directly on one or more of these foundations, providing deeper, contribution-specific background as needed.

Part I.

Analog Function Approximation

Part Overview

Flexible Electronics (FE) operate under severe physical constraints that make digital computation expensive. In the flexible IGZO platform targeted throughout this thesis, the dominant source of area and power consumption in a sensor-processing pipeline is the analog-to-digital interface: digitizing raw sensor signals at adequate resolution can consume the majority of the system’s transistor budget [6]. One attractive mitigation is to move part of the computation into the analog domain, *before* the ADC, thereby reducing or entirely avoiding digitization overhead for functions that can be approximated with sufficient accuracy in continuous-time circuitry.

Analog function approximation pursues this goal systematically. Rather than implementing a mathematical function in a conventional digital processor, the function is realized directly in the analog domain using elementary circuit primitives[43]. The approach is well-matched to FE for three reasons. First, sensor signals are inherently analog, so no conversion is needed to enter the processing chain [157]. Second, IGZO circuits operate at low frequencies (a few hertz to a few kilohertz), which relaxes bandwidth requirements and makes slow but area-efficient analog topologies competitive [145]. Third, the availability of a well-characterized process design kit (PDK)—such as PragmatIC’s FlexIC process—enables systematic circuit-level simulation of analog building blocks that can then be composed into more complex computational structures [125].

This part presents two main contributions in the domain of analog function approximation. Chapter 3 introduces the *Analog Building Block* (ABB) framework and demonstrates its use for approximating arbitrary univariate functions via second-order Bézier splines in a fully analog KAN architecture, establishing the hardware substrate and quantifying the fundamental area–accuracy trade-off. Chapter 4 extends this to a hardware–software co-design methodology—the Analog Kolmogorov-Arnold Network (AKAN)—that incorporates circuit-level error models and hardware-aware pruning directly into the KAN training loop, substantially reducing area and power while maintaining approximation accuracy.

Together, these two contributions demonstrate a clear progression: from a proof-of-concept analog implementation of a KAN to a principled co-design framework that closes the loop between physical circuit behavior and model training. This co-design perspective—jointly optimizing computational structure and hardware substrate—is a central theme that recurs across all parts of this thesis.

3. Function Approximation Using Analog Building Blocks in Flexible Electronics

3.1. Motivation and Problem Statement

Wearable sensors based on FE must process analog signals—such as biological markers or environmental measurements—in real time with minimal energy expenditure [103, 85]. Tasks such as computing a running average of heart rate, evaluating a threshold-based classifier, or applying a calibration curve to a raw sensor reading all require the evaluation of smooth nonlinear functions [54]. In digital implementations, each function evaluation requires full digitization followed by arithmetic processing, with the ADC alone often consuming the majority of system power [6, 94].

Analog function approximation offers a fundamentally different approach: the function is evaluated continuously in the electrical domain, consuming no conversion energy and operating at the signal’s native frequency [109]. The challenge is that FE analog circuits are resource-constrained—only n-type transistors are available, device matching is limited by process variability, and total transistor budgets are low—so approximation must be achieved with a small, fixed set of reusable circuit primitives[99]. Any proposed solution must therefore satisfy three simultaneous requirements: it must be expressible using only the available N-type transistor technology; it must be compositional, so that complex functions can be built from simpler verified sub-circuits; and it must be parameterizable, so that different target functions can be realized by changing only passive component values rather than redesigning the topology.

Kolmogorov-Arnold Networks (KANs) provide a structured mathematical framework for decomposing a multivariate function into a sum of univariate spline functions, each of which can be approximated independently with modest circuit resources [92, 83, 82]. Specifically, the KAN representation theorem guarantees that any continuous multivariate function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ can be expressed as:

$$f(x_1, \dots, x_n) = \sum_{q=1}^{2n+1} \Phi_q \left(\sum_{p=1}^n \phi_{q,p}(x_p) \right), \quad (3.1)$$

where $\phi_{q,p}$ and Φ_q are learned univariate functions realized as splines. This decomposition is ideally suited to analog implementation: each univariate spline is an independent circuit module, and the outer summation is a passive resistor network.

This chapter presents the first fully analog implementation of a KAN in FE, built from a library of five Analog Building Blocks (ABBs), and evaluates its area and power trade-offs against a digital reference.

3.2. Analog Building Block Design

Five ABB types are required to implement second-order Bézier splines in the analog domain: inversion, subtraction, addition, multiplication, and squaring. Each block is designed using only N-type transistors, resistors, and capacitors compatible with the Pragmatic FlexIC process, and operates with input and output ranges of $[-0.5\text{ V}, 0.5\text{ V}]$ [94]. All blocks share a common supply of $V_{dd} = 1.0\text{ V}$ and $V_{ss} = -1.0\text{ V}$.

The design rationale for each block is as follows:

Inversion. A two-transistor, one-resistor circuit makes $A' = -A$ through resistor-transistor logic. This is the smallest and most frequently reused building block, serving as an internal subroutine for the subtraction and addition blocks. The inversion is inherently linear and introduces negligible non-ideality.

Subtraction. The identity $A - B = A + (-B)$ is exploited: inversion is first applied to B , and the result $-B$ is combined with A using a two-transistor, one-resistor, one-bias-voltage circuit. The bias voltage V_{bias} extends the effective operating range of the subtractor and compensates for the quiescent operating point shift introduced by the N-type-only topology.

Addition. Leveraging $A + B = A - (-B)$, addition is implemented by first inverting B and then applying the subtraction block. No additional transistors are required beyond those of the inversion and subtraction primitives, making this the most resource-efficient of the two arithmetic blocks.

Multiplication. The Gilbert cell topology provides the most compact voltage-domain multiplier realization compatible with the FlexIC process [41, 147, 116]. The design is adapted to use only N-type transistors, and a final subtraction stage of $V_{dd} - A$ corrects the output range back into $[-0.5\text{ V}, 0.5\text{ V}]$, ensuring compatibility with downstream blocks.

Squaring. Squaring exploits the fact that the input range is known *a priori*, enabling a simplified variant of the multiplier: both inputs of the Gilbert cell are tied together, and the final output range correction uses a fixed bias voltage V_{bias} instead of a full subtraction block. This reduces transistor count and static power while maintaining the required output range.

The schematics of all five ABBs are shown in Figure 3.1, and their hardware characteristics—transistor count, simulated area, and quiescent power—are summarized in Table 3.1.




Figure 3.1.: Schematics of the five Analog Building Blocks: (a) inversion, (b) subtraction, (c) addition, (d) multiplication, and (e) squaring. All circuits use only N-type transistors and resistors compatible with the PragmatIC FlexIC IGZO process.

Table 3.1.: Hardware characteristics of the five Analog Building Blocks in the PragmatIC FlexIC process ($V_{dd} = 1.0\text{ V}$, $T = 27^\circ\text{C}$). Area and power are obtained from post-layout Cadence Spectre simulations.

ABB	Transistors	Resistors	Area (μm^2)	Power (μW)
Inversion	2	1	524.14	21.45
Subtraction	4	2	835.24	0.115
Addition	6	3	1957.84	21.6
Multiplication	8	2	18578	98.4
Squaring	6	1	11539	50.12
Total (one spline)	—	—	72750	238.5

* Transistor counts reflect the ABB in isolation; shared inversion sub-circuits within addition are counted once per instance.

3.3. Spline Implementation

3.3.1. Bézier Spline Formulation

A second-order Bézier spline with control points P_0 , P_1 , P_2 evaluates as:

$$\mathbf{B}(x) = P_0 + (P_1 - P_0) 2x + (P_0 - 2P_1 + P_2) x^2, \quad (3.2)$$

which decomposes naturally into three additive sub-expressions, each involving at most one power of x :

- P_0 : a constant offset, provided as a fixed bias voltage.
- $(P_1 - P_0) 2x$: a linear term, realized with a single multiplication ABB.
- $(P_0 - 2P_1 + P_2) x^2$: a quadratic term, realized with the squaring ABB followed by a multiplication ABB.

The three sub-expressions are summed using two addition ABBs, yielding a complete spline circuit with two multiplication operations, one squaring operation, and two addition operations—five ABB instances in total [94]. The architecture of the spline implementation is illustrated in Figure 3.2.




Figure 3.2.: Architecture of the spline implementation from ABBs. The three additive sub-blocks (P_0) , $(P_1 - P_0)2x$, and $(P_0 - 2P_1 + P_2)x^2$ are summed to make a second-order Bézier spline.

3.3.2. Control Point Fitting

Given a target function $f : [-0.5, 0.5] \rightarrow \mathbb{R}$, the segment is partitioned into S equal sub-intervals, and a separate Bézier spline is fitted to each segment. For a single segment spanning $[x_a, x_b]$, the normalized input $t = (x - x_a)/(x_b - x_a) \in [0, 1]$ is substituted into Equation (3.2). The three control point values P_0, P_1, P_2 are then determined by solving a least-squares problem:

$$\{P_0^*, P_1^*, P_2^*\} = \arg \min_{P_0, P_1, P_2} \sum_{j=1}^M [f(x_j) - B(t_j)]^2, \quad (3.3)$$

where M sample points $\{x_j\}$ are drawn uniformly from the segment. This fitting is performed entirely in software before the circuit is instantiated; the hardware only implements the evaluation of the fitted polynomial, with P_0, P_1, P_2 encoded as passive bias voltages and resistor ratios.

3.4. KAN Assembly

Multiple splines are combined into a KAN by applying weighted accumulation across all spline outputs. In the analog domain, this is achieved by connecting a series resistor to the output of each spline and tying all resistor outputs to a common node [154]. By Millman's theorem, the resulting voltage at the common node is a weighted average of the individual spline outputs, with weights inversely proportional to the corresponding resistor values [99]. Uniform weighting is achieved with equal-valued resistors; non-uniform weighting requires resistor sizing inversely proportional to the desired weights [155]. A bias resistor R_{bias} connected to a reference voltage is included to set the quiescent operating point.

The architecture of the full analog KAN is shown in Figure 3.3. The construction is entirely passive at the aggregation layer—no additional active devices are introduced by the MAC operation—which keeps the transistor count proportional to the number of splines rather than growing super-linearly.




Figure 3.3.: Architecture of the analog KAN implementation. Multiple splines $B_0(x), \dots, B_i(x)$ are connected in parallel through resistors $R_0, \dots, R_i$ to realize a weighted sum at the common output node $\text{KAN}(x)$. A bias resistor R_{bias} sets the quiescent output level.

3.5. Experimental Evaluation

3.5.1. Simulation Setup

All simulations are performed in Cadence Spectre with the PragmatIC FlexIC PDK (second-generation Helvellyn 2.1.0), at $V_{dd} = 1.0$ V, $V_{ss} = -1.0$ V, and 27°C . A ramp input signal spanning -0.5 V to 0.5 V over 10 ms with a $1\ \mu\text{s}$ timestep is used as the test stimulus. Area is estimated from transistor and passive component counts using the FlexIC design rules, and power is extracted as average DC dissipation across the full ramp sweep. On Figure 3.4, we present the layout of the single spline.




Figure 3.4.: Layout view of a single spline.

3.5.2. Spline Approximation Accuracy

Figure 3.5 shows the spline approximation quality for three representative target functions drawn from sensor-processing use cases: a sigmoid activation (relevant to neural classification), a square-root calibration curve (relevant to sensor nonlinearities), and a sinusoidal segment (relevant to periodic biosignal processing). In all three cases, the hardware-implemented spline closely tracks the ideal curve, with peak absolute errors remaining below 7.58%.




Figure 3.5.: Spline approximation accuracy for two representative target functions. Top row: ideal (dashed) vs. hardware-implemented (solid) spline output. The corresponding Normalized Mean Percentage Errors (NMPE) are included for both cases.

3.5.3. Comparison Against Digital Baseline

The analog KAN is compared against a digital implementation of the same second-order Bézier spline at 8-bit precision synthesized using the FlexIC digital standard-cell library. Results are summarized in Table 3.2.

Table 3.2.: Comparison of the analog KAN against digital spline implementations in the PragmatIC FlexIC process. Area and power are normalized to the 8-bit digital reference.

Implementation	Area (norm.)	Power (norm.)	Max error (%)
Digital, 8-bit	1.00×	1.00×	< 0.5
Analog KAN (ours)	0.008×	0.894×	7.58

Normalized area: analog KAN achieves a 125× reduction relative to the 8-bit digital baseline. Power reduction is 10.59%.

The 125× area reduction comes at the cost of a bounded approximation error of up to 7.58%, arising from both design-level non-idealities in the ABB implementations and parasitic coupling between cascaded circuit stages. For the target wearable sensing applications—where errors of this magnitude are typically tolerated after end-to-end system calibration—this trade-off is acceptable. It is worth noting that the 8-bit digital reference does not include ADC cost, which would further increase the digital implementation’s area and power when operating on a raw analog sensor signal [49, 150].

3.6. Summary and Contributions

This chapter has presented the first fully analog KAN implementation in flexible IGZO electronics, built from a library of five reusable Analog Building Blocks. The key contributions are:

1. A complete library of five ABBs—inversion, subtraction, addition, multiplication, and squaring—designed exclusively for the N-type IGZO PragmatIC FlexIC process, each verified by Cadence Spectre simulation against the FlexIC PDK.
2. A systematic methodology for fitting second-order Bézier control points to arbitrary target functions via least squares, composing ABBs into splines, and assembling splines into a full analog KAN via a passive resistor network.
3. A 125× area reduction and 10.59% power reduction compared to an 8-bit digital spline implementation, at the cost of a bounded approximation error of up to 7.58%.

The ABB library and spline composition methodology established here serve as the hardware substrate for the hardware–software co-design framework presented in Chapter 4, where the same circuits are revisited under a training-aware pruning strategy that substantially reduces area and power while maintaining approximation accuracy.

4. Hardware–Software Co-Design of Kolmogorov-Arnold Networks for Efficient Analog Function Approximation

4.1. Motivation

The ABB-based KAN presented in Chapter 3 demonstrates that analog function approximation is feasible in FE and offers compelling area savings over digital alternatives. However, that work treats the hardware as a fixed implementation target: spline parameters are chosen in software to best fit the target function, and the hardware is designed to realize those parameters as faithfully as possible. Two important limitations arise from this decoupling.

First, the analog hardware introduces systematic approximation errors that depend non-trivially on the specific values of the spline coefficients k_0, k_1, k_2 . These errors arise from non-ideal transistor behavior, parasitic resistance and capacitance, and supply-voltage coupling within the N-type-only Gilbert cell topology. Because the training procedure is unaware of this hardware behavior, the fitted model does not account for the actual approximation quality it will achieve in silicon. In the worst case, a training procedure that selects coefficient values producing large hardware deviations will outperform a hardware-naïve model in simulation while performing significantly worse on the physical circuit.

Second, not all spline coefficients contribute equally to the overall approximation quality across the input domain [140]. Some terms—particularly the constant offset k_0 or the quadratic term k_2 for nearly-linear functions—have negligible impact on the output and can be removed to save area and power without meaningful accuracy loss. A principled pruning strategy that evaluates hardware cost simultaneously with approximation error is therefore essential.

This chapter presents the **Analog Kolmogorov-Arnold Network (AKAN)**, which addresses both limitations through a hardware–software co-design methodology. AKAN incorporates a simulation-based circuit-level error model into the KAN training process and employs a two-level pruning strategy at both the spline-coefficient and KAN-network levels, yielding hardware implementations that are simultaneously more area-efficient and more accurate than the baseline approach [99].

4.2. Proposed Co-Design Methodology

The proposed AKAN framework, illustrated in Figure 4.1, consists of four interconnected stages: hardware architecture characterization, error modeling, pruning-aware KAN training, and performance evaluation. The stages are designed to be iterable: after an initial training pass, the pruning step may reveal that further training with a simplified architecture recovers accuracy, motivating a second training round.




Figure 4.1.: Block diagram of the proposed AKAN hardware–software co-design methodology. SPICE simulations characterize hardware splines and generate error models; KAN training incorporates these models alongside hardware costs; pruning reduces complexity at spline and network levels.

4.2.1. Hardware Architecture and Cost Characterization

The AKAN hardware adopts the same ABB-based spline architecture introduced in Chapter 3. The three spline parameters k_0 , k_1 , k_2 are derived from the Bézier control points via the transformations:

$$k_0 = P_0, \quad k_1 = (P_1 - P_0) \times 2, \quad k_2 = P_0 - 2P_1 + P_2, \quad (4.1)$$

and correspond directly to bias currents and resistor sizing ratios in the analog implementation [99].

To characterize the relationship between coefficient values, hardware cost, and approximation fidelity, a parametric sweep of $10 \times 10 \times 10 = 1,000$ coefficient configurations is simulated in Cadence Spectre, with each coefficient taking one of ten discrete values uniformly sampled from $[-0.5 \text{ V}, 0.5 \text{ V}]$. For each configuration, a ramp input from -0.5 V to 0.5 V is applied over 10 ms with a $1 \mu\text{s}$ timestep, yielding the hardware-implemented spline output.

4.2.2. Error Modeling

For each of the 1,000 simulated coefficient configurations, the deviation between the hardware-implemented spline output $\hat{\mathbf{B}}(x; k_0, k_1, k_2)$ and the ideal mathematical spline $\mathbf{B}(x; k_0, k_1, k_2)$ is characterized across the full input range. The error at input sample x_j is defined as:

$$\epsilon(x_j; k_0, k_1, k_2) = \hat{\mathbf{B}}(x_j; k_0, k_1, k_2) - \mathbf{B}(x_j; k_0, k_1, k_2), \quad (4.2)$$

and the root-mean-square error (RMSE) across the input range is used as the scalar summary metric for training:

$$\mathcal{E}(k_0, k_1, k_2) = \sqrt{\frac{1}{M} \sum_{j=1}^M \epsilon(x_j; k_0, k_1, k_2)^2}. \quad (4.3)$$

The resulting error function $\mathcal{E}(k_0, k_1, k_2)$ —characterizing how hardware non-idealities vary with specific parameter values—is integrated into the KAN training process as an additive regularization term. Unlike prior approaches that assumed a fixed error contribution independent of parameter values [99], the AKAN error model captures the input-dependent variation of approximation accuracy, producing training targets that are consistent with what the physical circuit will compute.

The validity of the error model is assessed by comparing predicted RMSE values against fresh Spectre simulations at coefficient configurations not included in the training sweep.

4.2.3. Pruning Strategy

Pruning in AKAN operates at two distinct levels.

Spline-level pruning. Individual polynomial coefficients (k_0 , k_1 , or k_2) are set to zero based on their contribution to overall approximation accuracy. Setting $k_0 = 0$ removes the constant offset sub-block, simplifying the spline to a linear-plus-quadratic form; setting $k_1 = 0$ removes the linear term, yielding a constant-plus-quadratic form; setting $k_2 = 0$ yields a purely linear approximation. Each simplification directly reduces the number

of ABB instances instantiated, as illustrated in Figure 4.2. The area and power savings associated with each pruning configuration are determined directly from the hardware cost map established in Section 4.2.1.




Figure 4.2.: Spline architectures before and after pruning. (a) Original second-order Bézier spline. After setting individual coefficients to zero: (b) $k_0 = 0$, (c) $k_1 = 0$, (d) $k_2 = 0$. Each pruning step removes one or more ABB instances, yielding a direct reduction in area and power.

KAN-level pruning. The impact of spline-coefficient pruning on the overall approximation accuracy is assessed across the complete input space for each spline in the network. Given that a KAN may contain multiple splines with different coefficient configurations, a pruning decision at the individual spline level can interact with approximation quality at other nodes. To prevent local simplifications from compounding into unacceptable system-level degradation, only pruning configurations that maintain the end-to-end approximation error within a specified bound $\epsilon_{\max}$ are retained. The search over valid pruning configurations proceeds greedily: splines are ranked by their potential area saving, and each candidate pruning is accepted only if the system-level RMSE constraint is satisfied.

4.2.4. Hardware-Aware KAN Training

The KAN is trained on application-relevant benchmark functions with the circuit error model incorporated as an additive regularization term in the loss function:

$$\mathcal{L} = \underbrace{\frac{1}{N} \sum_{i=1}^N [f(x_i) - \text{KAN}(x_i)]^2}_{\text{approximation loss}} + \lambda \underbrace{\sum_s \mathcal{E}(k_0^{(s)}, k_1^{(s)}, k_2^{(s)})}_{\text{hardware regularization}}, \quad (4.4)$$

where the sum over s indexes all splines in the network, and λ is a weighting hyperparameter that controls the relative importance of hardware fidelity versus mathematical accuracy. The regularization term penalizes coefficient configurations that produce large hardware deviations, implicitly guiding the optimizer towards regions of the coefficient space where the physical circuit more faithfully realizes the intended spline.

After training, the two-level pruning strategy of Section 4.2.3 is applied to identify and remove non-essential spline coefficients. The model is then retrained from the pruned architecture (warm-started from the pre-pruning weights) to recover any accuracy degraded by the structural simplification. The full training pipeline is summarized in Algorithm 1.

4.3. Experimental Results

4.3.1. Experimental Setup

Experiments are conducted on the PragmatIC FlexIC PDK (Helvellyn 2.1.0), using Cadence Spectre for all analog simulations. Benchmarks are drawn from four datasets representative of wearable sensing applications [51, 130, 32, 67]: the Iris classification dataset (sigmoid and softmax activations), PPG-DaLiA (photoplethysmography signal features), ECG5000 (electrocardiogram features), and the Individual Household Electric Power Consumption dataset (power calibration curves). For each dataset, a set of target functions is extracted from the learned activation shapes of a small reference neural network trained on the dataset. All AKAN configurations are evaluated on three metrics: approximation RMSE relative to the ideal mathematical function, total silicon area (computed from the hardware cost map), and total quiescent power (similarly from the cost map). The regularization weight λ is selected via a coarse grid search on a validation split, and the system-level error bound $\epsilon_{\max}$ is set to the value that allows at most 1% absolute accuracy degradation relative to the unpruned baseline.

Algorithm 1 AKAN Hardware-Aware Training and Pruning

Require: Target function f , hardware error map $\mathcal{E}$, hardware cost map (A, P) , error bound $\epsilon_{\max}$, regularization weight λ

- 1: Initialize KAN with full second-order Bézier splines (all k_0, k_1, k_2 free)
 - 2: **Phase 1 — Hardware-aware training:**
 - 3: **for** each training epoch **do**
 - 4: Compute approximation loss $\mathcal{L}_{\text{approx}}$ on mini-batch
 - 5: Look up hardware regularization $\mathcal{L}_{\text{hw}} = \lambda \sum_s \mathcal{E}(k^{(s)})$
 - 6: Update parameters via gradient descent on $\mathcal{L} = \mathcal{L}_{\text{approx}} + \mathcal{L}_{\text{hw}}$
 - 7: **end for**
 - 8: **Phase 2 — Spline-level pruning:**
 - 9: **for** each spline s , ordered by potential area saving **do**
 - 10: **for** each candidate pruning $\{k_0 = 0\} \cup \{k_1 = 0\} \cup \{k_2 = 0\}$ **do**
 - 11: Apply pruning candidate; evaluate system-level RMSE
 - 12: **if** RMSE $\leq \epsilon_{\max}$ **then**
 - 13: Accept pruning; update hardware cost map entry for spline s
 - 14: **break**
 - 15: **end if**
 - 16: **end for**
 - 17: **end for**
 - 18: **Phase 3 — Retraining:**
 - 19: Retrain pruned KAN (warm start) with $\mathcal{L}_{\text{approx}}$ only
- Ensure:** Pruned AKAN parameters; area, power, and RMSE for performance reporting
-

4.3.2. PVT Robustness Analysis

To assess the practical reliability of the AKAN under IGZO-relevant non-idealities, we conducted a comprehensive PVT characterization across 27 operating conditions: 3 process corners (SS, TT, FF) $\times$ 3 supply voltages (0.9 V, 1.0 V, 1.1 V) $\times$ 3 temperatures (0°C, 27°C, 40°C), covering the full wearable healthcare operating range [86, 64].




Figure 4.3.: System-level robustness analysis across 27 PVT corners ($V_{DD} \in \{0.9, 1.0, 1.1\}$ V, Temp. $\in \{0, 27, 40\}^{\circ}\text{C}$, and FF/TT/SS process corners). The top row illustrates NMPE stability, demonstrating hardware-aware compensation, while the bottom row shows the total power consumption in log-scale.

The system-level robustness of the AKAN spline across 27 distinct PVT operating conditions is summarized in Figure 4.3. As observed in the top panels, the NMPE remains remarkably stable within the wearable temperature range ($27\text{--}40^{\circ}\text{C}$), with a maximum deviation (ΔNMPE) of only 0.74 pp relative to nominal conditions. Even at the 0°C extreme, the worst-case NMPE is contained at 2.50%, confirming that the KAN training framework effectively compensates for corner-induced shifts by adapting the spline weights. These results indicate that the derived error model captures the dominant systematic hardware behavior across the evaluated environmental conditions.

Regarding power stability (Fig. 4.3, bottom), the AKAN exhibits the expected leakage increase at elevated temperatures—a characteristic of unipolar IGZO TFTs. However, the nominal power consumption (TT, 1.0 V, 27°C) remains highly efficient at $241.8\text{ }\mu\text{W}$. By leveraging the proposed coefficient-level pruning, this power footprint can be further reduced by up to 50% depending on the target function, maintaining a competitive power-to-performance ratio suitable for long-term wearable monitoring.

4.3.3. Area and Power Results

Figure 4.4 shows the Pareto front of approximation RMSE versus Area $\times$ Power for the full AKAN framework and the unpruned AKAN baseline across all four benchmark datasets.

Each point on the AKAN front corresponds to a different choice of $\epsilon_{\max}$, tracing the achievable trade-off between hardware efficiency and approximation accuracy.




Figure 4.4.: Pareto distribution of NMPE versus AreaPower for the Iris Petal Length dataset approximating $\sin(x)$. The plot shows the baseline (3 nonpruned splines, star marker) and all possible pruning scenarios across the three splines in the network. Pruning configurations are labeled according to which coefficient (k0, k1, or k2) is removed from each spline. Multiple configurations achieve lower NMPE and reduced $\text{Area} \times \text{Power}$ compared to the baseline, demonstrating favorable Pareto improvement through strategic parameter removal.

For each dataset, the pruning strategy was applied across three spline types and 20 unique pruning combinations, with averages reported in Table 4.1. The results show that pruning frequently improves the approximation quality rather than merely preserving it. This behavior arises because removing hardware-expensive coefficients can also eliminate regions of the parameter space associated with larger circuit-level non-idealities, allowing the retrained AKAN to converge towards coefficient configurations that are simultaneously more hardware-efficient and more faithful to the target function.

Table 4.1.: Summary of pruning impact on average and best NMPE across datasets, together with average area and power reduction.

Dataset	Avg. NMPE (Base) ¹	Best NMPE (AP) ²	Avg. NMPE (AP) ³	Δ NMPE ⁴	Avg. AR(%) ⁵	Avg. PR(%) ⁶
PPG-DaLiA						
PPG Chest	0.26	0.01	-0.22	0.04		
PPG Wrist	0.19	-0.01	-0.24	-0.05	29.8	30.5
Acc X	-0.91	-0.05	-1.35	-0.4		
Acc Y	1.64	0.04	1.20	0.44		
Iris						
Sepal L.	3.08	-0.08	-2.46	0.62		
Sepal W.	4.90	-0.05	-1.36	3.24		
Petal L.	3.77	-0.07	-1.89	1.88	29.8	30.5
Petal W.	4.77	-0.06	-1.59	3.18		
ECG5000	1.73	0.04	1.24	0.49	29.8	30.5
IHEPC⁷	-2.37	-0.16	-4.70	-1.97	29.8	30.5
Average	1.7	-0.04	-1.09	0.75	29.8	30.5

¹ NMPE of the baseline configuration (non-pruned). ² Best NMPE obtained among all pruning configurations (AP = After Pruning). ³ Average NMPE across all pruning configurations.

⁴ Signed change in NMPE, where a positive value indicates that pruning moves the NMPE closer to zero, and a negative value indicates degradation. ⁵ Average Area Reduction (%). ⁶ Average Power Reduction (%). ⁷ IHEPC = Individual Household Electric Power Consumption dataset.

Across all four benchmarks, the pruned AKAN achieves up to a **5.5 $\times$ improvement in the Area $\times$ Power product** compared to the unpruned baseline, with average reductions of approximately 30%. The best-case result—55% area saving and 50% power saving—is achieved on the ECG5000 dataset, where several splines approximate nearly-linear functions for which the quadratic k_2 coefficient contributes negligibly to accuracy. Importantly, the pruning process does not simply reduce hardware complexity while maintaining acceptable approximation quality; in several cases, it also improves NMPE relative to the baseline configuration. This indicates that hardware-aware pruning can act as an implicit regularization mechanism, removing coefficient configurations that amplify circuit-level non-idealities and guiding retraining towards more robust operating regions of the hardware parameter space.

4.3.4. Comparison to Chapter 3 Baseline

To make the progression explicit, Table 4.2 compares the full AKAN framework against the unco-designed ABB KAN from Chapter 3 on a common benchmark. The AKAN improves approximation accuracy by incorporating hardware error feedback during training, and further reduces area and power through principled pruning, demonstrating that hardware-aware co-design is not merely beneficial but necessary for effective analog computing in resource-constrained FE.

Table 4.2.: Comparison of the pruned AKAN framework against the baseline ABB KAN from Chapter 3 on a common sigmoid benchmark. All metrics are normalized to the 8-bit digital spline reference.

Implementation	Area (norm.)	Power (norm.)	RMSE (%)
Digital spline, 8-bit	1.00×	1.00×	< 0.5
ABB KAN (Chapter 3)	0.008×	0.894×	7.58
AKAN pruned (ours)	$\approx 0.005\times$	$\approx 0.62\times$	< 7.0

The comparison highlights the transition from a hardware-naive analog implementation to a fully hardware-aware co-design methodology. While the ABB KAN already demonstrated substantial area reduction relative to digital implementations, the AKAN further improves the Area $\times$ Power trade-off through coefficient-aware pruning and retraining. By integrating circuit-level behavior directly into the optimization loop, the AKAN achieves lower hardware cost while maintaining or improving approximation fidelity relative to the original ABB-based approach.

4.4. Summary and Contributions

This chapter has presented the AKAN, a hardware–software co-design framework that jointly optimizes analog KAN implementations for area, power, and approximation accuracy in FE. The key contributions are:

1. A simulation-based circuit-level error model that captures hardware-induced approximation deviations as a function of spline parameter values (k_0, k_1, k_2), enabling training to anticipate physical circuit behavior rather than treating the hardware as a post-hoc implementation target.
2. A hardware-aware training loss (Equation 4.4) that incorporates the coefficient-dependent error model as a regularization term, guiding optimization towards co-efficient configurations that are both accurate and hardware-friendly.

3. A two-level pruning strategy that removes non-essential spline coefficients at both the per-spline and network levels, guided by hardware cost models and system-level accuracy constraints.
4. Up to 55% area and 50% power reduction compared to the unpruned baseline, with average improvements of approximately 30% across four benchmark datasets, and up to a 5 $\times$ improvement in the Area $\times$ Power product.

The AKAN framework establishes that hardware–software co-design is not merely beneficial but necessary for effective analog computing in FE: training a model without knowledge of the hardware it will execute on leads to suboptimal implementations that could be substantially improved by closing the loop between simulation and training. This principle—explicitly informing model optimization with hardware constraints and behavior—underpins all subsequent contributions in this thesis, where the same co-design philosophy is applied to ADC architectures, fault-tolerant systems, and end-to-end wearable classifiers.

Part II.

**ADC Co-design and
Co-optimization**

Part Overview

The analog-to-digital converter is the most area- and power-intensive component in virtually every FE classification system. In a representative printed MLP for stress monitoring, the ADC circuitry alone consumes between 60% and 90% of the total transistor budget, as illustrated in Figure 4.5 [6]. This observation motivates a fundamentally different design stance: rather than treating the ADC as a fixed component whose resolution and architecture are chosen a priori, the ADC should be co-designed with the downstream classifier, allowing each to adapt to the requirements and constraints of the other.

This part presents two contributions that progressively develop this co-design principle. Chapter 5 proposes the Binary Search ADC architecture for FE classifiers and introduces a high-level in-training pruning methodology based on a multi-objective genetic algorithm that generates bespoke, per-input ADCs with the minimal representation necessary for each sensor channel. Chapter 6 extends this to non-uniform quantization, replacing the genetic algorithm with a swarm-intelligence optimizer and incorporating data-driven threshold placement at the circuit level through a non-uniform resistor ladder; post-layout simulation under PVT conditions validates the approach on a suite of real biosignal datasets, achieving up to $12.5\times$ power savings and $35.9\times$ area reduction compared to uniform quantization.

Collectively, these contributions demonstrate that the ADC is not a peripheral concern but a primary co-design variable: optimizing it jointly with the classifier can recover the majority of the area and power that would otherwise be wasted on quantization precision that the classifier does not need.




Figure 4.5.: Transistor count breakdown of a printed MLP classification system for six UCI benchmark datasets. The ADC front-end (blue) dominates the total transistor budget in every case, confirming that ADC co-design is the primary lever for area and power reduction in FE classifiers [100].

5. Design and In-Training Optimization of Binary Search ADC for Flexible Classifiers

5.1. Motivation

Among the ADC architectures available for FE, the Binary Search ADC offers the most favorable combination of area, power consumption, and implementation simplicity, as summarized in Table 5.1. Flash ADCs achieve the highest conversion speed but require $2^N - 1$ comparators and an N -bit priority encoder, making them prohibitively large at resolutions above 2 bits [84, 74]. SAR ADCs achieve lower area at the cost of a DAC feedback loop and digital control logic, which introduce routing complexity incompatible with the limited metal layers available in FE platforms [124, 8]. Delta-Sigma ADCs reach high resolution through oversampling but require operational amplifiers with substantial gain-bandwidth, which cannot be realized efficiently with IGZO N-type-only transistors [59, 153]. The Binary Search ADC avoids these limitations: it performs N sequential comparisons rather than $2^N - 1$ simultaneous ones, produces a binary code directly without an encoder, and operates entirely from resistors and N-type transistors, making it uniquely compatible with the PragmatIC FlexICs fabrication platform [125].

Despite these advantages, existing Binary Search ADC designs for FE target generic, fixed-resolution interfaces without considering the specific requirements of the downstream classifier [19]. This chapter addresses both limitations jointly: it proposes an FE-specific Binary Search ADC architecture that reduces area compared to the prior state of the art, and it introduces an in-training pruning optimization that tailors the per-input ADC to each feature of an MLP classifier [100].

Table 5.1.: Qualitative comparison of ADC architectures for FE (scale 1–4, where 1 is the lowest).

Architecture	Area	Power	Speed	Application
Flash	4	4	4	High speed and resolution
SAR	3	3	3	Low power and medium resolution
Sigma-Delta	2	1	1	Low power and high resolution
Binary Search	1	2	2	Low power and easy implementation

5.2. Proposed Binary Search ADC Architecture

5.2.1. Design constraints of the FE technology

The PragmatIC FlexICs platform provides only N-type IGZO thin-film transistors, a high-value integrated resistor component, and four metal routing layers on a polyimide substrate. The absence of P-type transistors eliminates complementary logic styles and forces all analog and digital functions to be implemented with unipolar circuits [125]. In conventional Binary Search ADC designs based on CMOS technology, the switching network responsible for selecting the reference voltage at each comparison stage is typically implemented with transmission gates, which require both N-type and P-type transistors [88, 19]. For the FE context, the switching network must be re-architected using only N-type pass transistors, which introduces a voltage-drop limitation: an N-type pass transistor cannot pass a signal above $V_{DD} - V_{th}$ without a bootstrapping circuit. The proposed design handles this constraint through careful transistor sizing: transistors M0 and M1 within each comparator are sized identically to guarantee that V_{in} and V_{ref} are treated symmetrically, while M2 and M3 are matched to ensure balanced loading of the comparison result and its complement. This matching strategy exploits a structural advantage of the N-type-only topology: all transistors track the same process corners together, so differential mismatch between P-type and N-type devices—a common source of systematic comparator offset in CMOS designs—is absent by construction.

5.2.2. Comparator design

The core comparator is shown in Figure 5.1(c) and comprises 7 transistors and 2 resistors. Transistors M0 and M1 form the differential input pair, accepting V_{in} and V_{ref} at their gates. When V_{DD} is applied, M0 and M1 pull down their respective drain nodes according

to their gate overdrive voltages; the resistors connected from each drain to V_{DD} form the load, and the resulting voltage difference is amplified through M2 and M3. Transistors M4, M5, and M6 implement a cross-coupled latch that forces the two outputs, V_{out} and V'_{out} , to complementary rail voltages (0 or V_{DD}), thereby delivering a digital decision without a separate inverter stage. Producing both the comparison result and its logical complement simultaneously is a key architectural choice: it eliminates the explicit inverter inserted after every comparator in the baseline Binary Search ADC design [19], and it allows the control block to select the correct reference voltage for the next stage without any additional logic depth.

5.2.3. Three-bit Binary Search ADC

The full 3-bit Binary Search ADC is shown in Figure 5.1. The complete design comprises 5 comparators, 2 NOT gates, and a control block implemented with 9 transistors (T0–T8), for a total that is $2\times$ lower in area than the baseline Binary Search ADC of the same resolution and $5.4\times$ lower than a 3-bit Flash ADC [100].




Figure 5.1.: Proposed 3-bit Binary Search ADC [100]. (a) Block diagram with three cascaded comparison stages producing bits A2 (MSB), A1, and A0 (LSB). (b) Full circuit schematic: 5 comparators (COM0–COM4), a threshold amplifier (TA), and the 9-transistor control block. (c) Comparator schematic using N-type-only transistors, delivering both the comparison result and its complement. (d) Reference voltage waveforms at each comparison stage: constant midpoint at Stage 1, a two-level square wave at Stage 2, and a four-level staircase at Stage 3.

Operational principle. Stage 1 compares V_{in} against $V_{ref0} = V_{ref}/2$, the midpoint of the full input range, producing the Most Significant Bit (MSB), A2. The output A2 and its complement $\bar{A}2$ directly gate transistors T0 and T1, routing power to exactly one of the two comparators (COM1 or COM2) that will be active in Stage 2; the inactive comparator receives no supply, eliminating its static current. Stage 2 determines A1 by comparing V_{in} against either $V_{ref} \cdot 3/4$ (if A2 = 1) or $V_{ref} \cdot 1/4$ (if A2 = 0). Before Stage 3, transistor TA amplifies the Stage 2 output when A2 = 1, driving the gate of the next comparator to a full V_{DD} level; this amplification step is necessary because the drain voltage of COM0 at

the threshold condition equals $V_{DD} - V_{th}$, which is insufficient to fully switch the comparators at higher reference voltages. Stage 3 resolves the Least Significant Bit (LSB), A0, by comparing V_{in} against one of four possible reference levels ($V_{ref} \cdot k/8$, $k \in \{1, 3, 5, 7\}$) depending on the binary decisions A2 and A1. Transistors T2–T5 implement the multiplexing of these four levels in the control block, while T6 and T7 select the Stage 3 reference based on whether $A2 > V_{ref}$.

Area comparison. Table 5.2 compares the proposed design against the baseline Binary Search ADC [19] and a 3-bit Flash ADC. For the 3-bit case, the proposed design occupies $17,679 \mu\text{m}^2$ compared to $35,722 \mu\text{m}^2$ for the baseline Binary Search ADC and $95,066 \mu\text{m}^2$ for the Flash ADC—a $2\times$ and $5.4\times$ reduction respectively. The power consumption of the proposed design (360 nW) is essentially identical to the baseline (365 nW), confirming that the area savings come entirely from removing the AND gate switching network, not from any power–accuracy trade-off. For the 4-bit case, the area advantage grows to $1.7\times$ over the baseline and $4.3\times$ over Flash, while power rises to 541.8 nW due to the additional comparator stages. It is important to note that the Flash ADC area reported includes the resistor ladder, comparators, and the digital priority encoder; the encoder alone accounts for approximately 9.8% of the Flash ADC area and 9.8% of its power—a cost entirely absent from the Binary Search ADC, which produces a naturally binary output code.

Table 5.2.: Full ADC comparison: area (μm^2) and power (nW) for 3- and 4-bit designs simulated at $V_{DD} = 1 \text{ V}$, $V_{ref} = 1 \text{ V}$, $T = 27^\circ\text{C}$, $f_{in} = 5 \text{ Hz}$ [100].

Architecture	Config.	Area (μm^2)	Power (nW)
Flash	3-bit	95,066	993.2
Flash	4-bit	212,635	2684
Binary Search (baseline)	3-bit	35,722	365.1
Binary Search (baseline)	4-bit	86,556	829.5
Binary Search (proposed)	3-bit	17,679	360
Binary Search (proposed)	4-bit	50,027	541.8

5.3. In-Training Binary Search ADC Pruning

5.3.1. Motivation and key observation

In a standard 3-bit Binary Search ADC, all $2^3 = 8$ quantization levels are available as possible outputs. For a given input feature of a trained MLP classifier, however, many of these levels may contribute nothing to the final classification decision: if a feature’s relevant information is concentrated in a narrow sub-range of its input domain, the comparisons that resolve levels outside that sub-range are always consistent and never change the digit that the classifier actually uses [14, 13, 12, 50, 3]. Removing those comparisons eliminates

the corresponding transistors, resistors, and control logic without affecting the classification output. The resulting “partial” or “pruned” ADC is bespoke to the specific (feature, dataset, classifier) triplet, and cannot be derived from the trained MLP alone—it requires a joint optimization that simultaneously determines the set of retained quantization levels and the MLP weights trained on those levels [145].

A subtle but important insight is that pruning can sometimes *improve* accuracy rather than merely preserve it. This occurs when the full-resolution ADC forces the MLP to fit data regions that are poorly represented in the training set [6], while the pruned ADC—by selecting levels aligned with the input distribution—implicitly regularizes the classifier, reducing overfitting. This effect is visible in the Pareto plots of Figure 5.2: pruned partial ADCs consistently achieve higher accuracy than their full-resolution counterparts at the same normalized area, because the quantization-aware training performed on the pruned ADC configuration effectively constrains the MLP to use representations that are well-supported by the data.

5.3.2. Pruning methodology: NSGA-II-driven exploration

The pruning procedure is framed as a multi-objective optimization problem with two competing objectives:

1. Maximize classification accuracy on the validation set.
2. Minimize the total ADC area, expressed in transistor count, using the high-level area model described below.

The search space is the set of all possible per-input ADC configurations: for each of the d input features, the pruning mask is a binary vector of length $2^N - 1$ (one bit per quantization level, indicating keep or prune) [119]. The combined search space has dimension $d \cdot (2^N - 1)$, which for $d = 13$ features and $N = 3$ bits yields 2^{91} candidate configurations—far too large for exhaustive evaluation.

The Non-dominated Sorting Genetic Algorithm II (NSGA-II) [39] is used to explore this space efficiently. NSGA-II is implemented via the open-source `PyMOO` framework [26] with binary string encoding: each individual encodes the full pruning mask across all input features together with the decimal point position of the fixed-point MLP coefficients, enabling joint optimization of the ADC representation and the classifier’s quantization-aware training (QAT) in a single evolutionary loop [108]. Mutation and crossover operators are applied with rates of 0.2% and 0.7%, respectively, and the population evolves over multiple generations; the non-dominated front from the final generation constitutes the Pareto set of ADC configurations.

High-level area model. Evaluating the ADC area of each candidate requires mapping from the pruning mask to a transistor count without running a full layout. A high-level area model is used, based on the design rules derived from the proposed Binary Search

ADC architecture: removing a quantization level a removes the transistor in the control block that routes V_{ref} for that level; if the pruning occurs in $V_{ref}/2$, the entire first-order comparator and half the tree are removed; if an entire level within the third comparison stage is pruned, the corresponding AND gate in the control block is also removed; if all values at a given level are pruned, the comparator for that stage is also removed. This model maps each pruning mask to a transistor count in $O(d \cdot 2^N)$ time, enabling fitness evaluation to keep pace with the genetic algorithm iterations.

5.3.3. Quantization-aware training on pruned ADCs

After each candidate ADC configuration is decoded, the MLP is trained from scratch using quantization-aware training (QAT) on the pruned representation. Inputs are quantized to the retained levels of the bespoke ADC before being passed to the MLP, and MLP weights are encoded as 8-bit fixed-point values with the decimal point position determined by the NSGA-II encoding. Training is performed with 70%/30% stratified train/test splits on each dataset, with balanced class distributions. MLP architectures follow those reported in prior work [115, 11] to enable fair comparison, and all designs are synthesized at a clock period compatible with the 5 kHz operating frequency typical of FE sensor systems [125].

5.3.4. Results

Pareto exploration. Figure 5.2 presents the Pareto fronts of accuracy versus normalized ADC area for 2-, 3-, and 4-bit designs across six datasets (BreastCancer, Cardio, Mammographic, RedWine, Seeds, WhiteWine). Area values are normalized to the corresponding Flash ADC baseline. The proposed Binary Search ADC (before pruning) already occupies the lower-left region relative to the Flash baseline, confirming the area advantage demonstrated in Table 5.2. Pruned partial ADCs push the Pareto front further toward the lower-left corner, achieving higher accuracy with less area than any full-resolution design at the same bit-width.




Figure 5.2.: Pareto fronts of classification accuracy vs. normalized ADC area for six UCI benchmark datasets, comparing 2-, 3-, and 4-bit designs [100]. Pruned Binary Search ADCs (*Ours*) consistently achieve higher accuracy with lower area than the corresponding full-resolution designs (*Baseline*), confirming that partial quantization eliminates representations that are harmful rather than merely redundant.

Quantitative summary. Table 5.3 summarizes the accuracy and transistor count for Flash ADCs, full Binary Search ADCs, and pruned Binary Search ADCs across 2-, 3-, and 4-bit resolutions. Replacing Flash ADCs with the proposed Binary Search ADC reduces the average transistor count by 46%, 48%, and 37% for 2-, 3-, and 4-bit designs, respectively, while increasing accuracy by approximately 3% on average due to the elimination of the encoder’s combinational delay. Applying the in-training pruning further reduces the transistor count by an average of 51% (relative to the full Binary Search ADC), for a compound reduction of 316–565% compared to the Flash ADC baseline, depending on resolution. The average accuracy *increases* by an additional margin after pruning, reaching 78% for both

3-bit and 4-bit designs compared to 77% for the un-pruned 3-bit baseline—consistent with the regularization effect described in Section 5.3.

Table 5.3.: Evaluation of printed MLP classification systems across ADC architectures and resolutions. Accuracy (%) is reported on the test set; area is in transistor count (TC) [100].

Res.	Accuracy (%)		Area (TC)			Gain (%)	
	Baseline	Pruned	Flash	Binary	Pruned	vs. Flash	vs. Binary
2-bit	73	78.2	423	235	134	57	316
3-bit	77	78.0	1138	523	249	48	458
4-bit	76	78.0	2676	981	474	37	565

Per-dataset pruning behavior. The degree of pruning varies substantially across datasets and input features. For Seeds (6 features, 3 classes), the area gains in transistor count range from $2.4\times$ (2-bit) to $8.4\times$ (3-bit) for up to 5% accuracy degradation. For WhiteWine (11 features, 7 classes), gains reach $24.2\times$ for 2-bit and up to $3.5\times$ for 4-bit at the same accuracy tolerance. This variability reflects genuine differences in feature informativeness: datasets with redundant or correlated features allow more aggressive pruning because multiple input channels can be reduced to near-binary resolution without compromising the classification boundary.

5.4. Summary and Contributions

1. A novel N-type-only Binary Search ADC architecture for FE, achieving $2\times$ area savings over the prior state of the art [19] and $5.4\times$ over Flash ADCs at 3-bit resolution, without any loss in power efficiency.
2. An NSGA-II-based in-training pruning methodology that jointly optimizes per-input ADC configurations and MLP classifier weights, reducing the transistor count by an average of 51% with respect to the full Binary Search ADC and improving classification accuracy by approximately 3% through implicit quantization-aware regularization.
3. A high-level area model that enables efficient fitness evaluation within the NSGA-II loop, making the exploration tractable for datasets with up to 11 input features and 4-bit resolution.

6. Low-Power Binary Search ADC with High-Level Non-Uniform Input Optimization

6.1. Motivation

The Binary Search ADC design and in-training pruning of Chapter 5 address the *number* of quantization levels per input feature. A complementary and equally important question is *where* to place those levels within the input range. Standard uniform quantization places levels at equal intervals, assigning the same resolution to every sub-range of the input domain regardless of how frequently that sub-range is observed in the data.

For biosignals—such as electrodermal activity (EDA) for stress detection, photoplethysmography (PPG) for heart-rate monitoring, or an accelerometer for balance assessment—the input distribution is fundamentally non-uniform [142, 22, 75]. EDA signals, for example, cluster near a baseline conductance level with rare high-amplitude excursions during stress events; PPG waveforms concentrate their information content in the systolic peaks and diastolic valleys, not in the intermediate flat segments. These distributions arise from underlying physiology and are therefore stable across measurement sessions and individuals within a target population, meaning that thresholds learned from a calibration set will generalize to deployment conditions. Uniform quantization wastes levels on low-density input regions and undersamples high-density regions, directly degrading classification accuracy.

A further motivation for non-uniform quantization in FE specifically is the static power structure of IGZO circuits: at a nominal supply voltage of 1.0 V, static power dominates (> 99% of total power consumption), because IGZO TFTs exhibit negligible dynamic switching current at the sub-kHz operating frequencies of wearable biosensor applications [145]. Consequently, the power cost of an ADC is determined almost entirely by the number of transistors it contains, not by their switching activity. Reducing the number

of quantization levels—by concentrating them in the high-density input regions and eliminating levels assigned to sparse regions—directly reduces transistor count and therefore power, without any dependence on the input signal’s dynamics [93].

This chapter presents a data-driven non-uniform quantization framework for the Binary Search ADC in a flexible ML healthcare system [101]. Unlike the uniform pruning of Chapter 5, which removes levels from the top of a fixed hierarchy, the non-uniform framework optimizes both the *number* and the *placement* of levels simultaneously, using a swarm-intelligence optimizer that searches the continuous space of threshold positions together with the discrete space of level counts. The framework is validated through post-layout simulation under Process, Voltage, and Temperature (PVT) variation using eight real biosignal datasets.

6.2. Non-Uniform Quantization Framework

6.2.1. Problem formulation

For input feature i , the quantization function is defined as:

$$Q_i(x) = c_{i,k} \quad \text{if } x \in [t_{i,k-1}, t_{i,k}), \quad (6.1)$$

where $k \in \{1, \dots, L_i\}$, $c_{i,k}$ is the k -th quantized centroid, and $\{t_{i,k}\}$ are the learned threshold positions. The set of centroids and thresholds for feature i is:

$$\mathcal{L}_i = \{c_{i,1}, \dots, c_{i,L_i}\}, \quad \#\mathcal{L}_i = L_i. \quad (6.2)$$

The design objective is the multi-objective optimization:

$$\min_{\{\mathcal{L}_i\}_{i=1}^N} \text{ADC_Area}(\{\mathcal{L}_i\}_{i=1}^N) \quad (6.3)$$

$$\max_{\{\mathcal{L}_i\}_{i=1}^N} \text{Accuracy}(\{\mathcal{L}_i\}_{i=1}^N), \quad (6.4)$$

where N is the number of input features, $\text{ADC_Area}(\cdot)$ models the total hardware cost given all active quantization levels, and $\text{Accuracy}(\cdot)$ measures the classification performance of the MLP trained on those levels. To explore both objectives in a unified way, a regularized fitness function is defined as:

$$\mathcal{F} = \text{Accuracy} - \lambda \cdot \sum_{i=1}^N \text{ADC_area}(i), \quad (6.5)$$

where λ is a trade-off parameter that controls the preference between accuracy and hardware efficiency. Larger λ values prioritize area savings at the expense of accuracy; smaller λ values favor classification performance. In practice, λ is tuned per dataset and hardware scenario, and the sensitivity of the Pareto front to λ provides insight into the inherent accuracy–area trade-off for each biosignal modality.

6.2.2. Artificial Bee Colony optimization

The search over $\{\mathcal{L}_i\}_{i=1}^N$ is conducted using the Artificial Bee Colony (ABC) algorithm [78, 77], a population-based swarm-intelligence method that maintains three types of agents: employed bees, onlooker bees, and scout bees.

Employed bees maintain candidate ADC configurations and explore their neighborhood by perturbing one randomly selected feature’s level count L_i ; if the new configuration improves the fitness $\mathcal{F}$, it replaces the previous one and the trial counter is reset. Onlooker bees probabilistically select high-fitness configurations from the employed bee population and apply the same perturbation logic, reinforcing exploitation of promising regions. Scout bees intervene when a configuration has not improved for a fixed number of trials—the abandonment threshold—replacing stagnant candidates with new random configurations to encourage exploration. This three-tier structure is illustrated in Figure 6.1.




Figure 6.1.: Artificial Bee Colony (ABC) optimization framework for non-uniform ADC configuration. Employed, onlooker, and scout bees explore ADC level combinations. Each candidate is evaluated via quantization, MLP inference, and fitness computation. The best-performing configuration is selected and mapped to ADC hardware.

ABC was chosen over the NSGA-II used in Chapter 5 for two reasons specific to the non-uniform setting. First, the threshold positions are continuous variables, making the binary-string encoding of NSGA-II less natural; ABC operates directly on real-valued level counts and uses K-Means internally to determine centroids, avoiding any discretization of the position space. Second, the objective function $\mathcal{F}$ is non-differentiable due to the discrete ADC area model and the non-smooth argmax of the classification accuracy; ABC handles this naturally without gradient information, whereas gradient-based refinement of NSGA-II individuals would require surrogate modeling. The fully parallel evaluation of bee populations further reduces wall-clock time: on the evaluated datasets, the average runtime of the ABC framework is 6 minutes per dataset, which is practical for the iterative design cycles typical of FE development.

6.2.3. Threshold determination via K-Means

For each candidate number of levels $\ell \in [L_{\min}, L_{\max}]$ for feature i , the framework applies K-Means clustering to the feature column $X[:, i]$ with ℓ clusters [105]. The cluster centroids C_i^ℓ become the quantized values, and the Voronoi boundaries between adjacent centroids become the decision thresholds T_i^ℓ . The dataset X is then quantized using these centroids and thresholds, and the fixed-point MLP is re-trained on the quantized data to evaluate the resulting accuracy $\mathcal{A}_i^\ell$.

The K-Means step is the link between the high-level optimization (which searches over level counts) and the circuit-level implementation (which requires specific threshold voltages) [93]. Its outputs directly determine the V_{ref} values that must be generated by the non-uniform resistor ladder, establishing a seamless path from data statistics to hardware parameters without any additional fitting step.

6.3. Four-Bit Non-Uniform Binary Search ADC Architecture

6.3.1. Circuit architecture

The proposed 4-bit non-uniform Binary Search ADC is shown in Figure 6.2. It comprises 10 comparators, 8 inverters, and 26 pass gates organized across four cascaded comparison stages, progressively narrowing the estimate of V_{in} by halving the search window at each stage. Stage 1 sets the MSB via comparator COM0. Stage 2 adjusts the reference using transistors T0 and T1 as pass gates, setting the second bit with COM2.0. Stage 3 further refines the result using pass gates T2–T7 and comparators COM3.0 and COM3.1. Stage 4 resolves the LSB using comparators COM4.0–COM4.3 and 12 transistors in the fourth-stage control block. Signal integrity across stages is maintained through double-inversion stages and additional transistor buffers, which prevent misclassifications due to weak signals at higher reference voltages.




Figure 6.2.: Architecture of the non-uniform 4-bit Binary Search ADC [101]. The non-uniform resistor ladder (left) generates data-distribution-adapted reference voltages. The central schematic realizes the four-stage hierarchical comparison with 10 comparators, 8 inverters, and 26 pass gates. Reference voltage profiles per comparison stage are shown on the right, illustrating how the search window narrows non-uniformly across the input range.

6.3.2. Non-uniform resistor ladder

Non-uniform quantization is realized at the circuit level by replacing the uniform resistor ladder—which divides V_{ref} into equal steps—with a non-uniform ladder whose resistance values are set to produce the learned threshold voltages at each tap. The reference voltage at node k of the non-uniform ladder is:

$$V_{ref,k} = \frac{\sum_{i=k+1}^N R_i}{\sum_{i=1}^N R_i} \cdot V_{ref}, \quad (6.6)$$

where R_i are the individual resistor values and N is the total number of resistor segments. This formulation shows that the voltage tap positions are fully determined by the ratios of consecutive resistances; changing the quantization thresholds reduces to changing resistance values, which does not require any additional comparators, inverters, or control logic beyond those already present in the uniform design. The total resistance of the ladder is fixed to match the uniform baseline, ensuring that the static current drawn by the ladder—and therefore its contribution to total static power—remains constant regardless of the threshold configuration.

Figure 6.3 illustrates the difference between uniform and non-uniform quantization for a representative 4-bit biosignal input, showing how non-uniform centroid placement allocates more quantization resolution to the high-density region of the input distribution.




Figure 6.3.: Comparison of 4-bit quantization schemes for a representative bimodal biosignal input [101]. **Top:** Uniform quantization divides the input range into 16 equal intervals regardless of data density, wasting resolution at the tails. **Bottom:** Non-uniform (data-driven) quantization places centroids near the two distribution peaks, preserving more classification-relevant information with the same number of levels.

6.3.3. Pruned non-uniform ADC

Once the ABC framework selects the optimal level count $\ell^* < 2^N$ for a feature, the corresponding ADC is pruned: unused branches of the Binary Search tree—comparators, pass transistors, and resistor ladder segments assigned to removed quantization levels—are simply not instantiated. Figure 6.4 shows an example pruned 4-bit Binary Search ADC for a feature that required only a subset of the full reference voltages. The unused branches are shaded to indicate that they are absent from the physical implementation. This tight coupling between the algorithmic threshold selection and the physical circuit implementation is enabled by FE’s low non-recurring engineering (NRE) cost: customizing one ADC per input feature in a silicon technology would be economically infeasible, but in FE, it is routinely achieved through mask-level customization without additional manufacturing steps.




Figure 6.4.: Example of a pruned 4-bit Binary Search ADC implementing non-uniform quantization for one biosensor channel [101]. Shaded components are absent from the physical implementation; only the branches corresponding to the retained quantization levels are instantiated, reducing both area and static power.

6.4. Post-Layout Validation Under PVT Conditions

6.4.1. Physical design

Layouts for both the 3-bit and 4-bit Binary Search ADC designs were created using the PragmatIC FlexICs Helvellyn 2.1.0 PDK and verified through Design Rule Check (DRC) and Layout versus Schematic (LVS) checks. Post-layout simulations were run after RC extraction using Calibre, capturing parasitic capacitances and resistances from the metal routing and device geometry. The physical layouts are shown in Figure 6.5.




Figure 6.5.: Physical layout of the proposed Binary Search ADC designs [101]. (a) 3-bit design: the three comparison stages, control block, and non-uniform resistor ladder are visible as distinct layout regions. (b) 4-bit design: the additional fourth comparison stage and second control block are included, increasing total area by approximately 3 $\times$.

6.4.2. Trade-off scenarios

The initial transistor sizing from Table 6.1 represents the baseline PVT operating point. Eight additional sizing scenarios (SC-A1 to SC-A4 and SC-P1 to SC-P4) are evaluated by modifying the bias resistance of early or late comparators, or by adjusting inverter resistances, as summarized in Table 6.2. The Figure of Merit (FoM) is defined as the normalized product of area and power; a lower FoM indicates a more efficient design. For 3-bit designs, area ranges from 0.047 mm^2 (SC-A1) to 0.063 mm^2 (SC-P3) and power ranges from 403 nW (SC-P3) to 500 nW (SC-A1), with SC-A3 consistently achieving the best FoM by biasing late comparators at high resistance. For 4-bit designs, SC-A3 and SC-P1 again produce the lowest FoM.

Table 6.1.: Simulation parameters used to evaluate robustness

Component	Element	Size
COM0&COM2.2	M0, M1	$W= 2\mu\text{m}$ $L= 600\text{nm}$
	M2 ,M3	$W= 4\mu\text{m}$ $L= 600\text{nm}$
	M4, M5, M6	$W= 8\mu\text{m}$ $L= 600\text{nm}$
	R	$r = 71M\Omega$
COM2.1¹	M0:M5	$W= 2\mu\text{m}$ $L= 600\text{nm}$
	R	$r = 71M\Omega$
COM2.0, COM3, COM4.0 to COM4.4	M0, M1	$W= 5\mu\text{m}$ $L= 600\text{nm}$
	M2 ,M3	$W= 2\mu\text{m}$ $L= 600\text{nm}$
	M4, M5, M6	$W= 2\mu\text{m}$ $L= 800\text{nm}$
	R	$r = 71M\Omega$
INV0:INV1 & INV3:INV8	M	$W= 3\mu\text{m}$ $L= 600\text{nm}$
	R	$r = 131.2M\Omega$
INV2	M	$W= 3\mu\text{m}$ $L= 600\text{nm}$
	R	$r = 11.4M\Omega$
T0:T21&TA0:TA3		$W= 2\mu\text{m}$ $L= 600\text{nm}$

¹COM2.1 only has Vout' as output, second R and M6 are removed

Table 6.2.: Comparison of 3-bit and 4-bit design trade-off scenarios; FoM is the normalized product of area and power [101].

Design	3-bit			4-bit		
	Area (mm ²)	Power (nW)	FoM	Area (mm ²)	Power (nW)	FoM
Baseline	0.056	442	1	0.152	581	1
SC-A1	0.047	500	0.96	0.147	624	1.04
SC-A2	0.052	449	0.95	0.149	592	1
SC-A3	0.049	432	0.87	0.136	639	0.98
SC-A4	0.048	457	0.9	0.156	659	1.16
SC-P1	0.061	411	1.03	0.157	564	1
SC-P2	0.055	421	0.94	0.179	568	1.15
SC-P3	0.063	403	1.03	0.182	542	1.11
SC-P4	0.059	416	1	0.152	579	1

6.4.3. Comparison against Flash and SAR ADCs

Figure 6.6 compares the proposed Binary Search ADC against Flash and SAR ADC designs at both 3-bit and 4-bit resolutions across all sizing scenarios. For 3-bit designs, Binary Search ADCs reduce area by $8.1\times$ – $10.8\times$ and power by $3\times$ – $3.8\times$ compared to Flash ADCs; SC-A3 shows the lowest Figure of Merit, balancing area and power most effectively. For 4-bit designs, area savings grow to $6.3\times$ – $8.5\times$ and power savings to $5.7\times$ – $6.9\times$ against Flash, with SC-A3 again showing the best FoM. Compared to SAR ADCs, the proposed Binary Search ADC consistently reduces area by up to $3.8\times$ (3-bit) and $1.7\times$ (4-bit), with power savings of $3.5\times$ for both cases. These comparisons account for all peripheral control logic, I/O interfaces, wiring, and ADC sub-blocks, not only the core comparator array.




Figure 6.6.: Area and power savings of the proposed Binary Search ADC relative to Flash and SAR ADC designs, for all trade-off scenarios at 3-bit (top) and 4-bit (bottom) resolution [101]. SC-A3 consistently achieves the best FoM, reducing area by up to 10.8 $\times$ and power by up to 3.8 $\times$ versus Flash.

6.4.4. PVT robustness

The simulation matrix covered supply voltage $V_{DD} = 1.0$ V nominal with $\pm 10\%$ variation (0.9 V and 1.1 V), temperature at 0 $^{\circ}$ C, 27 $^{\circ}$ C, and 85 $^{\circ}$ C, offset voltage $V_{off} \in \{-6.8$ mV, 0 mV, +6.8 mV $\}$, and transistor model corners (Typical, Fast, Slow). Input frequencies of 1, 5, 25, 500, 750, and 1000 Hz were evaluated to characterize conversion accuracy across the full biosignal frequency range.

The percentage deviation from the expected output code is computed as:

$$\text{Deviation} = 100 \times \frac{|A - B|}{|A| + X}, \quad X = \frac{V_{dd}}{32}, \quad (6.7)$$

where A is the expected value, B is the simulation result, and X accounts for the noise floor (half an LSB). Results, summarized in Figure 6.7, show that at nominal conditions ($T = 27^{\circ}$ C, $V_{off} = 0$, Typical model) and frequencies up to 500 Hz, deviation remains below 10%, confirming correct ADC operation across the full biosignal sensing bandwidth. Deviation grows toward 22.47% at 1 kHz under the Fast corner, reflecting the inherent speed–accuracy trade-off of the sequential comparison architecture. Temperature effects are the dominant source of deviation at 85 $^{\circ}$ C, indicating that thermal calibration—feasible through resistor-ladder re-programming—would extend reliable operation to higher temperatures. Supply voltage variation has a smaller impact: at $V_{DD} = 1.1$ V power rises to

684 nW (18% increase) while at $V_{DD} = 0.9$ V it drops to 484 nW (17% decrease), demonstrating that reducing supply voltage is an effective power management strategy for battery-constrained wearable deployments.




Figure 6.7.: Accuracy deviation of the proposed Binary Search ADC under PVT conditions [101]. Deviation remains below 10% for all temperature and offset combinations up to 500 Hz under nominal supply, confirming robustness across the wearable biosignal frequency range. The $\pm 10\%$ V_{DD} variation (bottom panel) causes proportional power changes but does not significantly affect conversion accuracy.

6.5. Results

6.5.1. Non-uniform ADC evaluation on biosignal datasets

The non-uniform quantization framework is evaluated on eight biosignal and physiological datasets from the UCI Machine Learning Repository [42]: Arrhythmia, Balance, Cardio, Mammographic, Parkinsons, Vertebral 2C, Vertebral 3C, and WESAD. Figure 6.8 presents the Pareto fronts for each dataset, showing the trade-off between classification accuracy and the total number of ADC levels (normalized to the 4-bit uniform baseline). Each point on the Pareto front represents a different non-uniform ADC configuration produced by the ABC optimizer; the uniform 4-bit ADC baseline (full $2^4 = 16$ levels) is shown as a

reference marker. The non-uniform configurations consistently achieve higher accuracy while using fewer levels than the uniform baseline—on average 4× fewer levels for the highest-accuracy solutions.




Figure 6.8.: Pareto fronts illustrating the trade-off between classification accuracy and ADC levels for eight biosignal datasets [101]. Non-uniform ADC configurations (*This Work*) consistently achieve higher accuracy with fewer levels than the uniform 4-bit baseline, demonstrating the benefit of data-driven threshold placement across diverse biosignal modalities.

6.5.2. Comprehensive comparison against the uniform baseline

Table 6.3.: Comprehensive evaluation of non-uniform vs. uniform 4-bit ADC designs across biosignal datasets. B^1 denotes the uniform 4-bit ADC baseline [101].

Dataset	Accuracy (%)			Power (μ W)			Area (mm^2)		
	B^1	Ours	Gain	B^1	Ours	Gain	B^1	Ours	Gain
Arrh.	62	67	+5%	842.5	67.2	12.5×	220.4	6.13	35.9×
Balance	91	92	+1%	12.20	1.68	7.3×	0.094	0.024	4.0×
Cardio	88	88	0%	41.83	5.20	8.0×	10.94	0.48	22.7×
Mamm.	86	85	−1%	12.20	1.48	8.2×	3.19	0.12	27.7×
Park.	81	93	+12%	41.83	5.77	7.3×	10.94	0.51	21.4×
Vert. 2C	90	92	+2%	13.94	1.54	9.0×	3.65	0.14	26.2×
Vert. 3C	83	88	+5%	13.94	1.56	8.9×	3.65	0.14	25.9×
WESAD	89	94	+5%	12.20	5.53	2.2×	3.19	0.49	6.5×

Table 6.3 presents detailed accuracy, power, and area results for the best non-uniform configuration per dataset, compared against the uniform 4-bit ADC baseline (denoted B^1). Accuracy improvements range from −1% (Mammographic, where the non-uniform threshold placement inadvertently misses a narrow but discriminative feature boundary) to +12% (Parkinsons). Power reductions range from 2.2×

 (WESAD) to 12.5× (Arrhythmia), and area reductions from 4.0× (Balance) to 35.9× (Arrhythmia). The particularly large savings for Arrhythmia arise from its high number of input channels (over 250), each of which would require a full 4-bit ADC instance under a uniform design; the ABC framework assigns many of these channels very few levels (often 2 or 3), dramatically reducing per-channel ADC complexity.

6.5.3. Comparison with state-of-the-art non-uniform quantization

Table 6.4 compares the non-uniform quantization approach of this work against the top-down pruning strategy of [100] (Chapter 5), in terms of accuracy gain and quantization level reduction relative to the uniform 4-bit baseline. The key distinction between the two approaches is directional: the in-training NSGA-II pruning of Chapter 5 removes levels from the top of a fixed uniform hierarchy, while the ABC-based framework of this chapter grows levels from the bottom up, identifying the most informative thresholds first. This bottom-up strategy is particularly effective for datasets where the most discriminative boundaries are not aligned with any sub-division of the uniform grid—precisely the case

for biosignals with physiologically shaped distributions. Our method achieves comparable or higher accuracy gains across all datasets (ranging from -1% to $+12\%$) while achieving competitive level reductions ($2.4\times$ – $5.5\times$ relative to the uniform baseline), confirming the advantage of jointly optimizing threshold positions and level counts.

Table 6.4.: Comparison of the non-uniform quantization approach of this work against the state-of-the-art pruning method [100]. Accuracy gain and level reduction are reported relative to the uniform 4-bit ADC baseline.

Dataset	Accuracy Gain (%)		Level Reduction ($\times$)	
	[100]	Ours	[100]	Ours
Arrhythmia	3	5	4.2	4.1
Balance	-4	1	3.8	2.4
Cardio	1	0	2.8	5.5
Mammographic	-2	-1	3.3	3.3
Parkinsons	12	12	3.5	4.6
Vertebral 2C	-5	2	3.1	4.0
Vertebral 3C	-14	5	2.7	4.0
WESAD	3	5	3.6	3.9

6.6. Additional Insights

Several observations from the experimental results deserve explicit discussion, as they have implications beyond what is reported in the individual papers.

Accuracy improvement through quantization. The fact that non-uniform quantization with fewer levels can *improve* accuracy—as observed for Arrhythmia, Vertebral, Parkinsons, and WESAD—is a manifestation of quantization as regularization. By restricting the input representation to a smaller codebook aligned with the data distribution, the MLP is forced to learn boundaries that generalize across the limited set of quantized values, effectively reducing the capacity of the input encoding and preventing the classifier from fitting noise in the continuous input domain. This effect is strongest for high-dimensional datasets with many features, where the MLP trained on the full-resolution uniform ADC can overfit individual quantization artifacts.

The Mammographic exception. Mammographic is the only dataset where accuracy decreases by 1% under non-uniform quantization. Inspection of the Mammographic feature distributions reveals a nearly uniform marginal distribution for several features, in contrast to the bimodal or skewed distributions of the other datasets. For such features, K-Means clustering does not find a more informative centroid placement than the uniform

grid, and the reduction in level count removes genuinely useful resolution. This observation supports the recommendation that the non-uniform framework should be applied with a minimum level count constraint $L_{\min} \geq 4$ for features with approximately uniform marginal distributions, to avoid over-pruning.

Arrhythmia as a paradigmatic case. The extreme savings achieved for Arrhythmia ($35.9\times$ area, $12.5\times$ power) arise from two compounding effects: the dataset has 250+ features, and many of those features are EEG-like signals with highly skewed distributions concentrated near zero. For such features, the ABC optimizer assigns $L_i = 2$ (a single threshold), reducing the corresponding ADC to a single comparator. The compound effect of simplifying hundreds of ADC instances simultaneously produces the large absolute savings observed; the per-feature saving is more modest ($2\times$ – $5\times$) and consistent with the other datasets.

Static power as the fundamental ADC constraint in FE. Since static power represents over 99% of total power in IGZO FE circuits at 1.0 V supply, the correct metric for optimizing FE ADCs is transistor count—not energy per conversion, which is the standard figure of merit for silicon ADCs. This has a direct architectural implication: circuit techniques that reduce dynamic power (e.g., clock gating, charge recycling, sub-threshold operation) provide negligible benefit in the FE context, while techniques that reduce the number of always-on transistors (pruning, non-uniform resistor ladders, N-type-only topologies) provide proportional static power reduction. The non-uniform Binary Search ADC exploits this property by removing resistor ladder segments and ADC branches for unpopulated input regions, reducing static current through those branches to zero.

6.7. Summary and Contributions

1. A data-driven non-uniform quantization framework for Binary Search ADCs in flexible ML healthcare systems, jointly optimizing quantization threshold positions and MLP classifier weights using an ABC swarm-intelligence optimizer with a regularized fitness function.
2. A 4-bit non-uniform Binary Search ADC architecture tailored for FE, implemented with a non-uniform resistor ladder requiring no additional comparators or control logic beyond those of the uniform design; the threshold voltages are determined entirely by resistance ratios.
3. Up to $12.5\times$ power reduction and $35.9\times$ area reduction compared to the uniform 4-bit ADC baseline, alongside up to 12% accuracy gains, validated through post-layout simulation under PVT conditions on eight biosignal datasets.
4. Comparative analysis against Flash and SAR ADCs demonstrating up to $10.8\times$ area and $3.8\times$ power savings at 3-bit resolution, and against the state-of-the-art

uniform pruning strategy of Chapter 5 demonstrating competitive or superior accuracy–efficiency trade-offs across all evaluated datasets.

Part III.

Reliability and Fault Tolerance

Part Overview

The manufacturing processes underlying FE produce defect densities substantially higher than those in mature silicon CMOS technologies. The combination of low-temperature photolithographic deposition of IGZO semiconductors, the absence of rigid packaging, and the exposure to mechanical stress during operation results in circuit-level faults that would be rare or manageable in silicon but represent a first-order design challenge in FE [117, 128, 17]. Analog and mixed-signal circuits—and ADCs in particular, which form the critical interface between the physical sensor world and the digital processing backend—are disproportionately vulnerable: faults that produce errors in the ADC output code propagate directly into downstream processing and cannot be recovered by fault-tolerant digital logic [134, 9].

A further complication is that conventional test strategies for analog circuits rely on ADC-based voltage monitoring [149, 118], yet an ADC used as a test observer is itself fabricated in the same variable process as the circuit under test, making it vulnerable to the same supply fluctuations and threshold shifts that degrade the circuit it is supposed to monitor—a phenomenon termed *test invalidation*. This fundamental limitation motivates the oscillator- and spectral-signature-based BIST strategies developed in this part.

This part presents three contributions that address reliability and fault tolerance for ADC and mixed-signal circuits deployed in structural testing and on-chip monitoring scenarios. Chapter 7 develops a hierarchical fault injection and sensitivity analysis framework for Binary Search ADCs in IGZO, deriving selective redundancy strategies that improve single-fault coverage from 60% to 92% and multi-fault coverage from 34% to 77.3% with only 4.2% area overhead. Chapter 8 presents a distributed delay-based BIST architecture built around a low-power IGZO ring oscillator and VCO, achieving $1100\times$ area and $5600\times$ power reduction over prior designs, with 93% individual-fault and 88% multi-fault defect coverage at 3% power overhead. Chapter 9 introduces a frequency-domain solution using ratiometric RMS spectral signatures that convert fault-induced bias deviations into compact spectral observables, maintaining fault separability across global PVT variation without any ADC in the test path.

Cross-cutting both the fault-tolerant design and the BIST approaches is the observation that fault sensitivity in analog circuits is highly non-uniform: protecting the most critical few nodes—whether through targeted redundancy or strategically placed monitoring oscillators—recovers the majority of fault coverage at a fraction of the cost of uniform protection. The application of these reliability techniques to power management circuits and

the Flash ADC interface, in the context of a complete wearable and implantable healthcare system, is presented in Part IV (Chapter 10).

7. Fault Tolerant Design of IGZO-based Binary Search ADCs

7.1. Motivation

While Chapter 10 addresses fault tolerance in Flash ADCs through comparator-level hardening, Binary Search ADCs present a fundamentally different fault propagation scenario. In a Flash ADC, a faulty comparator affects only its corresponding thermometer-code bit—an error localized to one position [74]. In a Binary Search ADC, a fault in an early conversion stage misdirects the entire subsequent search process: the MSB comparator’s decision in Stage 1 determines which half of the input range is searched in Stage 2, and any incorrect MSB corrupts all subsequent bits, producing catastrophic output code errors regardless of the correctness of later comparators [88]. This hierarchical fault propagation—absent in parallel topologies—demands a methodology specifically designed for the cascaded comparison architecture [96].

Furthermore, conventional fault simulation for analog circuits is both compute-intensive and poorly adapted to FE characteristics: existing CMOS models do not capture IGZO-specific failure modes such as large threshold-voltage spread, leakage-dominated static behavior, and mechanical-stress-induced opens [141, 160, 102, 10, 69]. A methodology that combines transistor-level defect characterization with system-level propagation assessment—while remaining tractable for exhaustive multi-fault evaluation—is needed.

7.2. Hierarchical Fault Injection Framework

7.2.1. Overview

The proposed framework follows a five-step hierarchical workflow, illustrated in Figure 7.1: (1) transistor-level component characterization; (2) fault library generation at subcircuit level; (3) system-level fault propagation assessment; (4) multi-fault scenario analysis; and (5) performance evaluation with design hardening.




Figure 7.1.: Hierarchical fault injection workflow [96]. Left: five-step process from transistor-level characterization to performance evaluation with design hardening. Right: system-level propagation decision tree (fault is catastrophic if $DNL \geq 1$ LSB, marginal if $0.5 \leq DNL < 1$ LSB, benign otherwise) and the ADC schematic with selective redundancy elements highlighted.

7.2.2. Defect Modeling

Structural defects in IGZO-TFT circuits are modeled by injecting resistive parasitics into Cadence Spectre transistor-level netlists [23]. Open-circuit defects are represented by high-impedance paths (250 M Ω) at transistor terminals; short-circuit defects are represented by low-impedance bridges (10 Ω) between circuit nodes. Each transistor terminal (gate, drain, source) is tested independently, giving five fault types per transistor: OC_G, OC_D, OC_S, SC_DG, SC_GS. The choice of 250 M Ω for opens and 10 Ω for shorts is deliberately conservative: real resistive opens and partial shorts in IGZO TFTs span a wide range, and these values represent the extremes that produce the most pronounced behavioral deviation [141, 102, 79].

7.2.3. Fault Library and System-Level Classification

Individual comparator subcircuits are extracted from the ADC hierarchy and subjected to exhaustive fault injection to construct a fault library that maps specific defect locations to their functional impact. Faulty models are then substituted into the complete 3-bit Binary Search ADC netlist and full analog-to-digital conversion sequences are simulated with a 5 Hz sinusoidal input spanning the full 1 V range. The Differential Nonlinearity (DNL) is computed across the full-scale input range; faults are classified as:

- **Catastrophic:** complete converter malfunction or output codes stuck at incorrect values ($DNL \geq 1$ LSB).
- **Marginal:** significant DNL degradation ($0.5 \leq DNL < 1$ LSB), conversion remains operational.
- **Benign:** minimal DNL impact ($DNL < 0.5$ LSB), functionally acceptable.

7.2.4. Multi-Fault Scenarios

Single-fault analysis identifies individual vulnerabilities but fails to capture defect clustering effects that occur in realistic manufacturing scenarios. Multi-fault test cases are constructed by strategically combining faults based on their individual classification: pairs of catastrophic faults (worst-case combinations); catastrophic-plus-marginal combinations (realistic clustered defects); and multiple marginal faults (accumulated degradation effects). Both baseline and fault-resilient designs undergo these multi-fault stress tests.

7.2.5. Automated Simulation Environment

Figure 7.2 shows the automated simulation infrastructure, managed entirely by a custom Python framework that coordinates netlist manipulation, parametric fault injection, batch execution through Cadence Spectre and Ocean, and metric extraction. The pipeline supports parallel execution of faulty netlists, substantially reducing simulation time, and maintains full traceability of results to enable rapid and reproducible design iterations. Across baseline and proposed designs, more than 220 distinct fault configurations are evaluated per architectural variant.




Figure 7.2.: Automated fault simulation infrastructure [96]. A Python framework coordinates netlist parsing, resistive fault injection, parallel Spectre execution, and DNL-based metric extraction, enabling evaluation of more than 220 fault configurations per design variant.

7.3. Fault Sensitivity Analysis

Figure 7.3 presents the fault sensitivity heat map under single-fault injection for the baseline 3-bit Binary Search ADC. Each cell represents a specific transistor terminal (row: fault type; column: transistor) and is colored red (catastrophic), yellow (marginal), or green (benign).




Figure 7.3.: Fault sensitivity heat map of the baseline Binary Search ADC under single-fault injection [96]. Red: catastrophic ($\text{DNL} \geq 1 \text{ LSB}$); yellow: marginal; green: benign. Fault vulnerability concentrates in early-stage elements (COM0, the control block's TA and T0–T1), while late-stage comparators COM3 and COM4 show 0% baseline coverage under multi-fault scenarios.

The heat map reveals three key observations. First, fault vulnerability is highly non-uniform: catastrophic failures cluster in a small subset of transistors, concentrated in the differential input pairs and the output stages of the early comparators COM0 and COM1. This motivates selective rather than uniform redundancy. Second, the hierarchical architecture amplifies early-stage faults: an open-circuit fault on the gate of M1 in COM0 does not produce a fixed output error—it causes the MSB to be inverted for all inputs, forcing the entire second-stage decision into the wrong half of the input range and corrupting all output codes simultaneously. Third, the control block faults (transistors T0–T7, TA) are predominantly marginal rather than catastrophic, because a partially conducting control transistor narrows the reference voltage window without eliminating it—producing a stretched but not incorrect decision boundary. This structural insight, visible directly in the heat map, directly guided the placement of redundant elements in the proposed fault-tolerant architectures.

7.4. Vulnerability-Aware Fault-Tolerant Design

Based on the baseline vulnerability analysis, two fault-tolerant architectures are proposed.

Selective Front-End Redundancy (SFR) implements selective redundancy at the high-sensitivity transistors identified by the heat map. Redundant elements are added at: input amplification stages (TA.1, TA.2), pass-gate stages (T0.1, T1.1), and the first inverter stage

(INV0.1). All redundant devices maintain the original sizing ($W = 2\ \mu\text{m}$, $L = 600\ \text{nm}$). The redundancy is placed exclusively in the elements that, when faulty, cascade their error through the maximum number of subsequent stages.

Extended Comparator-Level Redundancy (ECLR) extends SFR with additional output-stage redundancy and an extra comparator transistor M7 in COM0, COM3, and COM4. M7 is introduced in the differential pair's output path, creating a secondary output node that is compared against the primary output; a mismatch flags a potential fault. ECLR targets the comparators that contributed most to catastrophic outcomes in the single-fault analysis—specifically the residual vulnerabilities in COM0 and the complete vulnerability of COM3 and COM4 in the baseline (both at 0% multi-fault coverage).

Figure 7.4 compares the physical layouts of the baseline and ECLR designs at both the comparator and full-ADC level.




Figure 7.4.: Physical layout comparison [96]. (a) Baseline COM3 comparator. (b) Fault-resilient COM3 with extended comparator-level redundancy (ECLR). (c) Complete baseline 3-bit Binary Search ADC. (d) Complete fault-resilient ECLR design, with 4.2% area overhead ($58,615\ \mu\text{m}^2$ vs. $56,260\ \mu\text{m}^2$).

7.5. Results

7.5.1. System-Level Coverage

Table 7.1 summarizes system-level fault coverage for all three design variants. SFR achieves 88.9% single-fault and 73.2% multi-fault coverage; ECLR reaches 92% and 77.3% respectively, compared to the 60%/34% baseline. Both proposed designs incur only 4.2% area overhead and 6% power increase. For comparison, conventional Triple Modular Redundancy (TMR) would require approximately 200% area overhead to achieve comparable single-fault coverage [18]; the selective approach reaches similar resilience at 2–3% of that resource cost by targeting only the vulnerability-critical locations identified through hierarchical analysis.

Table 7.1.: Fault coverage, area, and power for all design variants. Coverage: percentage of fault configurations with DNL < 1 LSB. Data from [96].

Design	Single FC	Multi FC	Area (μm^2)	Power (nW)
Baseline	60%	34%	56 260	334.1
SFR	88.9%	73.2%	58 615 (+4.2%)	354 (+6.0%)
ECLR	92.0%	77.3%	58 615 (+4.2%)	353.8 (+5.9%)

7.5.2. Per-Component Analysis

Table 7.2 presents multi-fault coverage per circuit component, revealing which blocks benefit most from the proposed hardening strategy.

Table 7.2.: Multi-fault coverage by circuit component for all three design variants. Data from [96].

Component	Baseline	SFR	ECLR
Control Block	51.52%	51.52%	51.52%
COM0	54.55%	63.64%	72.73%
COM1	17.65%	82.35%	94.12%
COM2	35.00%	90.00%	90.00%
COM3	0.00%	100.00%	100.00%
COM4	0.00%	83.33%	100.00%

Several observations merit emphasis. The output comparators COM3 and COM4, completely vulnerable in the baseline (0% multi-fault coverage), are fully protected in the proposed designs (100% in ECLR). COM1 shows the most dramatic improvement: from 17.65% baseline to 94.12% in ECLR. This large gain arises because COM1 was the most common source of cascading faults—its output gates the enable signal for Stage 3, so any COM1 fault that produces an incorrect enable propagates to both COM3 and COM4 simultaneously. The control block remains at 51.52% across all variants, representing a fundamental bottleneck: control-block faults tend to produce DNL patterns just below the catastrophic threshold (marginal classification), which resist improvement from the output-stage redundancy strategy. Addressing the control block would require a distinct redundancy strategy, such as duplicating the pass-gate network T2–T7. The SFR-to-ECLR progression shows diminishing returns in COM2 and COM3 (both saturate) but continued improvement in COM0 (+9.09 percentage points), COM1 (+11.77), and COM4 (+16.67), suggesting that a design between SFR and ECLR that targets only COM4’s remaining gap would be the most area-efficient next step.

7.5.3. Multi-Fault Interaction Examples

The multi-fault analysis reveals concrete examples of how selective redundancy contains fault propagation. When combining COM1 M0 OC_G (marginal in single-fault) with COM4 M4 OC_D (catastrophic in single-fault), the unmodified design produces catastrophic DNL degradation; both SFR and ECLR reduce this to marginal, demonstrating that the architectural improvements prevent individual defect classifications from compounding. When pairing two individually catastrophic faults—COM3 M3 OC_S and COM4 M4 OC_D—the proposed designs again limit degradation to marginal, confirming that selective redundancy at output stages prevents worst-case fault combinations from causing complete converter malfunction. The proposed designs can also downgrade marginal fault accumulation to benign: combining COM3 M1 OC_D (marginal) with COM2 M2 SC_GS (marginal) yields marginal DNL in the baseline but benign classification in both SFR and ECLR, demonstrating that the redundancy not only prevents catastrophic failures but also improves tolerance to accumulated moderate degradations.

7.6. Summary and Contributions

1. A hierarchical fault injection and sensitivity analysis framework for Binary Search ADCs in IGZO FE, adapted to the cascaded conversion architecture and technology-specific failure modes, with a fully automated Python simulation pipeline.
2. A fault sensitivity heat map that reveals the asymmetric vulnerability profile of the Binary Search ADC: early-stage comparators are catastrophically sensitive

while late-stage comparators are fully unprotected under multi-fault injection in the baseline.

3. A selective redundancy design strategy targeting the most fault-sensitive components (SFR and ECLR), improving single-fault coverage from 60% to 92% and multi-fault coverage from 34% to 77.6%, with only 4.2% area and 6% power overhead—achieving comparable resilience to TMR at 2–3% of its resource cost.

8. Distributed Delay-Based BIST for Mixed-Signal Circuits in Flexible Electronics

8.1. Motivation

Post-manufacturing and in-field testing of FE mixed-signal circuits is complicated by the absence of rigid packaging, the limited number of accessible pins, and the high device variability that makes fixed parametric limits impractical to apply globally [17, 128, 33]. Conventional ATE-based approaches require physical connections to each test point, which is incompatible with the form factor and deployment scenarios of wearable and conformal FE systems. Built-In Self-Test provides an on-chip solution, but conventional mixed-signal BIST architectures for silicon—which rely on ADC-based voltage observation [47, 149, 90]—introduce significant area overhead and are sensitive to the same process variations that affect the circuit under test.

This chapter presents a dual-purpose ring oscillator (RO) and voltage-controlled oscillator (VCO) architecture that serves simultaneously as a functional timing reference and as a distributed delay-based BIST monitor [95]. The key insight is that delay variation is not a nuisance to be minimized but a signal to be measured: inter-stage delay differences in a faulty RO encode information about the location and severity of defects, and lightweight digital logic at each stage can read this information without any ADC or external reference clock.

8.2. Proposed RO–VCO Architecture

8.2.1. Ring Oscillator

The RO uses three cascaded inverting stages (M1–M3) with resistive pull-ups (R5–R7) and a biasing voltage divider (R0, R4). Series resistors R1–R3 ensure sufficient gain margin, and the combination of R8, R9, and capacitor C sets the oscillation frequency. A deliberate asymmetry is introduced by connecting the pull-up of Stage 2 to the voltage divider node

rather than directly to V_{DD} ; this breaks the DC-locking symmetry that would otherwise prevent oscillation in a perfectly balanced three-stage ring. Under PVT variation, the RO remains functional across all evaluated corners, achieving 38–122 kHz at Fast and 37–85 kHz at Slow process corners.

8.2.2. Voltage-Controlled Oscillator

The VCO extends the RO by adding a single control transistor M0 in the discharge path of Stage 1. The gate of M0 is driven by V_{CONTROL} , modulating the effective impedance of the first stage's discharge path and thereby controlling the oscillation frequency with a single additional transistor and no additional passive components. Lower V_{CONTROL} values yield faster discharge and higher oscillation frequency; higher values slow the discharge and lower the frequency. This achieves wide-range frequency tunability while introducing negligible area overhead.

8.2.3. Digital BIST Logic

The BIST logic, shown in Figure 8.1, is embedded at each RO stage. Adjacent stage outputs are processed by XNOR gates, followed by an OR gate and a D flip-flop for synchronization. The flip-flop uses the oscillator's own output as its clock, eliminating any requirement for an external reference clock. Under nominal, fault-free operation, the consistent phase relationship between adjacent stages produces a BIST output of logic 0. A defect that alters the delay of one stage breaks this phase relationship; the XNOR output then toggles and the flip-flop captures a logic 1, flagging the fault condition. Loss of oscillation entirely—e.g., due to a catastrophic open-circuit defect—is handled by default: a static output is treated as a defect condition.




Figure 8.1.: Schematics of (a) the proposed RO, (b) the VCO extending the RO with a single control transistor M_0 , (c) the digital BIST block using XNOR and flip-flop logic for stage-wise delay comparison, and (d) the fully integrated VCO+BIST with per-stage monitoring [95].




Figure 8.2.: Layout views of (a) the VCO, (b) the BIST digital logic, and (c) the integrated VCO+BIST design implemented in the PragmatIC FlexIC Gen3 process [95]. The BIST block occupies $4,585 \mu\text{m}^2$, adding 12% area relative to the VCO alone.

8.3. Experimental Setup

All simulations were performed using Cadence Spectre with device models from PragmatIC's third-generation FlexIC process. The RO operated at $V_{DD} = 0.9 \text{ V}$; the VCO control voltage was swept from 0.65 V to 0.85 V at $V_{DD} = 0.9 \text{ V}$ and from 1.05 V to 1.41 V

at $V_{DD} = 1.5$ V. Component sizing is given in Table 8.1. The BIST was designed using the Gen3 FlexIC standard-cell library.

Table 8.1.: Device properties for the RO, VCO, and BIST circuits. Data from [95].

Component	Element	Size
RO	R0	50 k Ω
	R1–R7	172.47 k Ω
	R8	120 k Ω
	R9	5.6 M Ω
	M1, M2	$W = 20 \mu\text{m}, L = 600 \text{ nm}$
	C	1 pF
VCO	Same as RO	—
	M0	$W = 20 \mu\text{m}, L = 600 \text{ nm}$
	C	2 pF
BIST	Std-cell based (Gen3 FlexIC)	

Defect injection targets the resistive pull-ups R1–R3 and R5–R7. Short-circuit (SC) defects model resistive bridging during fabrication: the nominal 172 k Ω is replaced by 10 k Ω . Open-circuit (OC) defects model broken interconnects: resistance increases to 10 M Ω . Defects are injected one at a time and in all pairwise simultaneous combinations, for a total of 180 simulation runs.

8.4. PVT Characterization

Figure 8.3 shows the RO frequency and power as a function of V_{DD} across all process-temperature corners. The counter-intuitive observation—frequency *decreasing* with increasing V_{DD} —is a direct consequence of the resistive pull-up topology: since the pull-up resistance is fixed, a higher supply voltage does not increase the available charging current, but it does increase the required voltage swing across the stage, prolonging the charging time. This inverse relationship between V_{DD} and frequency is a fundamental property of resistor-loaded CMOS inverters operating far from the current-saturation regime, and it must be accounted for when interpreting BIST signatures under supply-voltage variation.




Figure 8.3.: Effect of V_{DD} on oscillation frequency and power per process-temperature corner [95]. Frequency decreases with increasing V_{DD} due to the resistive pull-up topology; the Fast corner achieves up to 122 kHz ($32.8 \mu\text{W}$) and the Slow corner is most power-efficient ($5 \mu\text{W}$).

Table 8.2 summarizes the PVT frequency and power ranges across the three process corners.

Table 8.2.: PVT corner analysis results for the RO. Data from [95].

Transistor Process	Frequency Range (kHz)	Power Range (μW)
Fast (F)	38–122	7.1–32.8
Typical (T)	37–97	5.7–23.9
Slow (S)	37–85	5.0–17.5

Despite the $3\times$ frequency spread between Fast and Slow corners, the BIST detection principle remains valid across all corners because it relies on *relative* inter-stage delay rather than absolute frequency. A golden-device or power-on-calibration reference signature is established once per device; subsequent comparisons measure deviations from that device’s own baseline, making the approach inherently robust to the inter-device process spread.

8.5. Defect Detection and BIST Output

Figure 8.4 shows the BIST output waveforms under three conditions: nominal operation, OC defect detection, and SC defect cases.




Figure 8.4.: BIST output waveforms after full VCO+BIST integration [95]. Left: nominal operation (flag remains logic 0). Centre: OC defect detection at Stage 1 pull-up R5 (flags raised in Stages 1 and 3; Stage 2 flag analysis in Figure 8.5). Right: SC defect cases where the flag is not activated (inherently tolerant behavior explained below).

The BIST achieves **93% defect coverage for individual faults** and **88% for multiple simultaneous defects**. The 7% of individual defects that remain undetected are concentrated in a specific scenario: an SC defect in Stage 2. Because Stage 2 connects its pull-up to the voltage divider node rather than directly to V_{DD} —the deliberate asymmetry that prevents DC locking—a short-circuit bridging the pull-up still leaves the stage operational but at a shifted frequency. The resulting VCO output oscillates at a different frequency but with a consistent phase relationship between stages, so the XNOR-based comparator observes no inter-stage deviation and does not raise a flag. This specific undetected fault is, by design, a fault-tolerant scenario: the asymmetric topology makes Stage 2 inherently more resistant to this defect type.

Figure 8.5 shows the VCO output waveforms for the principal undetected fault scenarios, confirming the cancellation mechanism.




Figure 8.5.: Detailed VCO outputs for cases where the BIST fault flag is not activated [95]. The 2nd-stage SC fault is undetected because the asymmetric V_{DD} connection of the middle stage makes it inherently tolerant to this defect: the oscillation continues at the same frequency with a consistent inter-stage phase, producing no detectable phase deviation in the XNOR-based comparator.

For multiple simultaneous defects, the 88% coverage (vs. 93% for individual defects) reflects cases where two defects produce opposing delay effects: one defect accelerates a stage while another slows it, and the net inter-stage phase shift falls below the BIST detection threshold. This cancellation effect is an intrinsic limitation of single-threshold XNOR-based comparators; adaptive threshold schemes that vary the sensitivity based on the nominal operating corner could mitigate it.

8.5.1. Comparison with the State of the Art

Table 8.3 compares the proposed designs against prior RO and VCO implementations in FE. The proposed RO and VCO operate at 0.9 V supply—the lowest reported for IGZO ROs—while achieving competitive frequency ranges. The power reduction is $1100\times$ lower area and $5600\times$ lower power than prior IGZO RO designs that operated at 20–25 V.

Table 8.3.: Comparison of proposed ROs and VCOs with FE state of the art. Data from [95].

Reference	V_{DD}	Freq.	Stages	Power (mW)	Area (mm ²)
IGZO RO (this work)	0.9 V	40 kHz	3	0.0089	0.0041
IGZO VCO (this work)	0.9–1.5 V	54–135 kHz	3	0.0083–0.017	0.0041
VCO+BIST (this work)	0.9–1.5 V	55–133 kHz	3	0.0091–0.017	0.0235
IGZO VCO [80]	5 V	327–560 Hz	—	1.3	n.r.
SWCNT VCO [81]	5 V	0.85–2.05 kHz	5	≈0.4	≈17.5
IGZO RO [34]	20 V	1.82–6.51 MHz	7	≈43	≈0.13
IGZO RO [61]	25 V	810 kHz	11	≈50	≈0.208
IGZO RO [143]	25 V	406 k–2.1 MHz	5–7	≈13–43	≈1.6–4.5

8.6. Summary and Contributions

1. A dual-purpose IGZO RO/VCO architecture operating at 0.9 V supply, achieving 1100× area and 5600× power reduction over prior IGZO RO designs, suitable for deployment as a distributed BIST monitor. The asymmetric Stage-2 V_{DD} connection is identified as both the mechanism that prevents DC locking and the source of inherent tolerance to Stage-2 SC faults.
2. A distributed delay-based BIST that embeds per-stage lightweight digital detection logic, achieving 93% single-fault and 88% multi-fault defect coverage with only 3% power overhead. The approach uses the oscillator’s own output as the BIST clock, eliminating any external reference clock requirement.

9. Spectral Signatures for Parametric Fault Detection in Flexible Electronics

9.1. Motivation

The delay-based BIST of Chapter 8 detects structural defects through time-domain stage-delay monitoring. An orthogonal approach exploits the frequency-domain response of the circuit under test (CUT) as a spectral signature. A VCO converts the CUT's DC output bias—which deviates from its fault-free value under parametric fault conditions—into a frequency shift; the redistribution of spectral energy across harmonics that results from this shift is then extracted as a compact, pass/fail-ready test observable.

The fundamental problem with conventional ADC-based voltage monitoring in FE is test invalidation [97]. ADC accuracy in IGZO degrades under supply-voltage fluctuations characteristic of FE: matched comparators are highly sensitive to threshold-voltage spread; stable references require unavailable P-type devices; and pseudo-CMOS digital logic incurs large area and static power. Since the BIST circuitry is fabricated in the same variable process as the CUT, an ADC-based observer produces incorrect pass/fail decisions under PVT variation regardless of whether the CUT is actually faulty. The spectral signature approach avoids this entirely by using ratiometric normalization: absolute signal levels are irrelevant because only their ratio is compared against a threshold. Just as a differential amplifier rejects common-mode noise by computing the difference between two correlated signals, the ratiometric signature rejects common-mode PVT variation by dividing the HPF RMS by a stored fault-free reference—collapsing the $5\times$ inter-corner absolute spread into a $\leq 9.5\%$ normalized band.

9.2. Spectral Signature Methodology

9.2.1. Signal Flow

Figure 9.1 illustrates the complete signal flow of the proposed circuit. During test, the CUT is set to its nominal input stimulus. Its DC output bias — V_{CONTROL} — drives the VCO. A parametric fault in the CUT (threshold shift, mobility degradation, resistive deviation) produces a deviation ΔV in this DC bias, which the VCO translates into a frequency shift. The frequency shift redistributes spectral energy: the fundamental component decreases while higher harmonics increase (and vice versa, depending on fault polarity). Two first-order RC filters—a low-pass filter (LPF) and a high-pass filter (HPF)—extract RMS energy from complementary spectral bands [133]:

$$S_{\text{LPF}} = \text{RMS}(v_{\text{LPF}}), \quad S_{\text{HPF}} = \text{RMS}(v_{\text{HPF}}). \quad (9.1)$$

Their cutoff frequencies bracket the nominal VCO operating range:

$$f_{\text{LPF}} \in [0, f_{\text{VCO,low}}], \quad f_{\text{HPF}} \in [f_{\text{VCO,high}}, \infty). \quad (9.2)$$




Figure 9.1.: Complete signal flow of the proposed frequency-domain test signature circuit [97]. The CUT output bias drives the VCO; LPF and HPF extract spectral energy; normalized RMS signatures are compared against threshold τ to produce a binary pass/fail decision without any ADC or digital counter.

9.2.2. Ratiometric Normalization and Pass/Fail Decision

The normalized signature is defined as:

$$S_{\text{norm}} = \frac{S_{\text{HPF}}}{S_{\text{ref}}}, \quad (9.3)$$

where S_{ref} is the HPF RMS value measured at the fault-free operating point for each corner during a one-time post-manufacturing calibration step. Per-corner normalization eliminates the absolute PVT spread (up to $5\times$ between extreme corners), collapsing all corners into a narrow normalized band of $\leq 9.5\%$ in the high-sensitivity region. A fault is declared when:

$$\text{Decision} = \begin{cases} \text{FAIL} & \text{if } S_{\text{norm}} > \tau_{\text{high}} \text{ or } S_{\text{norm}} < \tau_{\text{low}}, \\ \text{PASS} & \text{otherwise,} \end{cases} \quad (9.4)$$

with thresholds $\tau_{\text{high}} = 1.10$ and $\tau_{\text{low}} = 0.90$. In the high-sensitivity region ($V_{\text{CONTROL}} < 1.8 \text{ V}$), a 300 mV fault produces a 22% change in S_{norm} , yielding a margin of $>12\%$ over the $\leq 9.5\%$ PVT spread—sufficient to set the thresholds without false positives. The LPF output provides a complementary low-sensitivity reference that reduces susceptibility to false positives; the HPF output maximizes sensitivity to fault-induced harmonic distortion.

9.2.3. Circuit Implementation

Figure 9.2 shows the full circuit implementation. The VCO uses the same ring-oscillator topology as Chapter 8, adapted for the 3 V supply of the Gen3 FlexIC process. No external clock, frequency counter, ADC, or high-speed digital logic is required. The LPF and HPF are realized as first-order RC networks whose cutoffs separate fundamental and harmonic content. The complete test signature is readable via two standard I/O pins by any microcontroller already present in the wearable system.




Figure 9.2.: Circuit-level implementation of the frequency-domain test signature circuit [97]. The IGZO-based VCO (M0–M4, R0–R10) drives first-order LPF and HPF filter networks. Component sizing: VCO R1–R7 = 172.47 k Ω , C = 0.3 pF; HPF R = 251.27 k Ω , C = 2 pF; LPF R = 172.47 k Ω , C = 2 pF.

9.3. Experimental Setup

All simulations were performed using Cadence Spectre with device models from PragmatiC's third-generation FlexIC process at $V_{DD} = 3\text{ V}$. The VCO control voltage was swept from 1.2 V to 2.7 V. Periodic steady-state (PSS) simulations were used to extract the steady oscillatory behavior; spectral components were then processed through the LPF and HPF. PVT variations were evaluated across 9 corners (Typical/Fast/Slow $\times$ T0/T1/T2) at three temperatures (0 °C, 27 °C, 85 °C). Results are normalized to the fault-free corner (T1_T) at $V_{\text{CONTROL}} = 1.2\text{ V}$.

Parametric faults are modelled as controlled additive deviations ΔV in V_{CONTROL} relative to the fault-free operating point V_0 , following established analog fault simulation practice [141]. Four fault magnitudes are evaluated: $|\Delta V| \in \{75, 150, 300, 450\}\text{ mV}$, corresponding to approximately 5%, 10%, 20%, and 30% of the full 1.2–2.7 V operating range. Each magnitude is applied at three representative operating points: $V_0 = 1.2\text{ V}$ (high-sensitivity), $V_0 = 1.5\text{ V}$ (intermediate), and $V_0 = 2.0\text{ V}$ (low-sensitivity). In IGZO TFTs, threshold-voltage shifts of 100–300 mV under bias stress and process spread are well-documented [122]; in a source-follower readout stage, such a shift translates directly into an equal DC output deviation, placing realistic faults well within the evaluated range.

9.4. Results

9.4.1. VCO Characteristics

Figure 9.3 shows the oscillation frequency as a function of V_{CONTROL} under fault-free PVT conditions. The VCO exhibits a monotonic tuning characteristic across the full range, with higher oscillation frequencies at lower control voltages (250–350 kHz range), consistent with the resistive-load topology. Although the frequency-voltage relationship is not strictly linear, its monotonic nature ensures an unambiguous correspondence between the control voltage and the oscillation regime—the only property required by the proposed methodology.




Figure 9.3.: Oscillation frequency of the proposed VCO as a function of V_{CONTROL} under fault-free conditions [97]. The monotonic characteristic across 250–350 kHz ensures an unambiguous mapping from CUT bias deviation to frequency shift, which is the foundational requirement of the spectral signature approach.

9.4.2. RMS Signatures Under Fault-Free PVT Conditions

Figure 9.4 (left panel) shows the RMS value of the second harmonic under nominal conditions as a function of V_{CONTROL} . The signature is monotonically increasing, rising from 0.14 at 1.2 V to 0.2 at 2.7 V. Sensitivity is non-uniform: in the lower region (1.2–1.5 V) the RMS increases by 22.5% (from 0.14 to 0.17) for a 300 mV shift, while in the higher region (2.0–2.7 V) the increase is only 5.2% (from 0.19 to 0.2) for a 700 mV shift. This non-uniformity directly explains the operating-point dependence of fault coverage: the same fault magnitude produces a much larger signature change in the low- V_{CONTROL} region, making that region intrinsically more suitable for fault detection.

The right panel of Figure 9.4 shows absolute RMS values across all nine PVT corners. At $V_{\text{CONTROL}} = 1.2$ V, values range from 0.04 (T2-F corner) to 0.22 (T0-F corner)—a spread exceeding 5 $\times$. Critically, however, the ordering with respect to V_{CONTROL} is preserved for all corners with no curve crossings. This property—monotonicity preserved across corners—is what makes normalization effective: once each corner is normalized to its own fault-free reference, the collapsed band is narrow enough to separate fault-induced deviations from PVT noise.




Figure 9.4.: RMS value of the second harmonic vs. V_{CONTROL} : nominal conditions (left) and all nine PVT corners (right) [97]. The absolute spread exceeds $5\times$ between extreme corners, but the ordering is preserved for all corners—the critical property enabling per-corner ratiometric normalization.

9.4.3. Normalized Signatures and Fault Coverage

Figure 9.5 presents the normalized signature S_{norm} under fault injection across all nine PVT corners at the three representative operating points. The dashed and dotted horizontal lines indicate the pass/fail thresholds $\tau_{\text{high}} = 1.10$ and $\tau_{\text{low}} = 0.90$.




Figure 9.5.: Normalized RMS signatures S_{norm} under fault injection for all nine PVT corners at three operating points [97]. The green region indicates the detectable fault zone ($S_{\text{norm}} > \tau_{\text{high}}$) and the red region the undetectable zone. At $V_0 = 1.2$ V, slow-process corners clearly exceed τ_{high} for $|\Delta V| \geq 150$ mV; at $V_0 = 2.0$ V, all corners remain within the pass band regardless of fault magnitude.

Table 9.1 summarizes the fault coverage (FC) across all nine PVT corners and both fault polarities, confirming the strong dependence on operating point.

Table 9.1.: Fault coverage vs. fault magnitude $|\Delta V|$ and operating point. Evaluated across 9 PVT corners $\times$ 2 fault polarities. Thresholds: $\tau_{\text{high}} = 1.10$, $\tau_{\text{low}} = 0.90$. Data from [97].

$ \Delta V $	FC at $V_0 = 1.2$ V	FC at $V_0 = 1.5$ V	FC at $V_0 = 2.0$ V
75 mV (5%)	44%	0%	0%
150 mV (10%)	65%	28%	0%
300 mV (20%)	69%	56%	6%
450 mV (30%)	100%	76%	17%

The incomplete coverage at small fault magnitudes (75 mV at $V_0 = 1.2$ V: 44%) is concentrated in fast-process corners, where higher absolute VCO gain paradoxically reduces relative sensitivity: the signature change from 75 mV at fast corners is only $\approx 4\%$, below τ_{high} . Slow-process corners are highly sensitive (S_{norm} reaches 2.37 for a +300 mV fault at $V_0 = 1.2$ V), confirming that the method is complementary: corners where the CUT is most at risk (slow IGZO devices with lower drive strength and higher susceptibility to threshold shift) are exactly the corners where the spectral signature is most sensitive.

9.4.4. Sensitivity Analysis

Figure 9.6 (right panel) shows the normalized RMS signatures across all PVT corners at $V_{DD} = 3\text{ V}$, and the left panel shows the derivative of the normalized signature with respect to V_{CONTROL} —i.e., the per-corner sensitivity. Maximum sensitivity occurs in the lower control-voltage region (1.2–1.5 V): at the T1-T corner, the normalized RMS increases from 1.0 to 1.225—a 22.5% change for a 300 mV shift. At higher voltages (2.0–2.7 V), the same 700 mV shift produces only 3.2% change.




Figure 9.6.: Sensitivity of the normalized RMS signature to V_{CONTROL} deviations (left) and normalized RMS signatures across all PVT corners at $V_{DD} = 3\text{ V}$ (right) [97]. Maximum sensitivity in the 1.2–1.5 V region yields 22.5% signature change for 300 mV faults vs. $\leq 9.5\%$ PVT spread, providing a $>12\%$ margin for reliable threshold placement.

9.4.5. Comparison with ADC-Based Monitoring

Table 9.2 compares the proposed test circuit against representative ADC-based monitoring solutions in the same IGZO-based FE technology. The proposed circuit achieves the smallest area (0.0064 mm^2) and lowest power (0.017 mW) across all compared solutions, representing a $26\times$ area reduction over the most compact ADC (Binary Search ADC at 3 V) and nearly $1500\times$ over the SAR implementation. Crucially, this comparison understates the true ADC overhead: ADC accuracy is itself sensitive to supply-voltage variations, comparator matching, and reference stability; and the ADC digital output still requires additional threshold comparison logic to produce a pass/fail decision, while the spectral signature circuit produces a pass/fail directly at its analog output.

Table 9.2.: Comparison with ADC-based flexible test circuits. Data from [97].

Reference	Resolution	Power (mW)	Area (mm^2)
This work	—	0.017	0.0064
Binary ADC [96]	3-bit	0.071	0.17
Flash ADC	4-bit	0.126	1.18
SAR ADC [8]	7.8-bit	2.93	9.4

The physical layout of the complete test circuit, shown in Figure 9.7, confirms the compact implementation occupying $6,390 \mu\text{m}^2$ with clearly delineated VCO, HPF, and LPF blocks.




Figure 9.7.: Physical layout of the proposed frequency-domain test signature circuit [97]. The VCO, HPF, and LPF blocks are visible; total area is $6,390 \mu\text{m}^2$. The compact footprint represents a $26\times$ area reduction versus the most compact ADC-based monitoring solution.

9.5. Summary and Contributions

1. A VCO-based frequency-domain test signature circuit for parametric fault detection in analog and AMS CUTs, requiring no ADCs, precision references, or CUT modification. The methodology avoids test invalidation by design: ratiometric normalization suppresses common-mode PVT variation in the same way a differential amplifier suppresses common-mode noise, collapsing the $5\times$ inter-corner absolute spread to a $\leq 9.5\%$ normalized band.
2. Definition and validation of ratiometric RMS spectral signatures using the second harmonic of the VCO output, exhibiting monotonic sensitivity to operating-point deviations (22.5% change per 300 mV fault in the high-sensitivity region) while preserving ordering across all nine PVT corners.

3. Implementation in the PragmatIC FlexIC Gen3 process, occupying $6,390\ \mu\text{m}^2$ and consuming $16.7\ \mu\text{W}$, achieving $26\times$ area reduction versus compact ADCs and nearly $1500\times$ versus SAR implementations. The signature is readable via two standard I/O pins, enabling in-field test without ATE.
4. An analysis identifying that fast-process corners exhibit reduced relative sensitivity (a paradoxical consequence of higher VCO gain), motivating per-corner adaptive thresholds as a natural extension.

Part IV.

Healthcare Applications and Clock Generation

Part Overview

The contributions presented in the preceding parts—analogue function approximation, ADC co-design, and reliability—are not ends in themselves but enabling technologies for a concrete class of applications: wearable and implantable healthcare monitoring systems built on flexible substrates. This part demonstrates the end-to-end integration of these technologies into complete flexible healthcare systems, addressing four complementary challenges: cross-layer reliability in the presence of physiological environments and mechanical deformation; systematic design-space exploration of the ML classification pipeline; analogue-domain feature extraction that bypasses the ADC bottleneck; and stable clock generation for synchronous AMS operation.

Stress detection from physiological signals provides the primary application context throughout, as this is a pervasive health concern linked to cardiovascular disease, sleep disorders, and cognitive impairment, yet existing monitoring solutions are rigid, silicon-based, and expensive—incompatible with the continuous, unobtrusive, and low-cost monitoring required for broad healthcare access [1, 164]. The combination of FE’s mechanical flexibility, sub-dollar fabrication cost, and ultra-low power consumption makes it an ideal platform for wearable stress monitoring patches, provided the system-level design challenges—reliable power delivery, fault-tolerant signal acquisition, classifier complexity, and clock stability—are addressed jointly.

Chapter 10 opens the part with a cross-layer reliability framework that pairs passive implantable sensing with robust wearable electronics: a gain-boosted IGZO low-dropout regulator (ReFlex-LDO) with integrated structural BIST, a hierarchical fault-tolerant Flash ADC design achieving up to 97.7% single-fault coverage, and a passive RF-based sensor that monitors hydrocephalus shunt patency without any on-board power. Chapter 11 presents the first comprehensive design space exploration of ML classifiers for stress monitoring in FE, evaluating over 1200 configurations across three ML algorithms, three feature selection methods, and four precision levels, and identifying Pareto-optimal designs achieving 94% accuracy at $9\ \mu\text{W}$. Chapter 12 introduces the first analogue feature extractors in FlexIC technology, replacing the dominant digital feature extraction bottleneck with a mixed-signal co-design where statistical features are extracted continuously in the analogue domain before a single shared SAR ADC; the resulting system achieves $48\times$ area and $70\times$ energy-per-inference reductions over purely digital co-design. Chapter 13 closes the part—and the thesis—by addressing the universally neglected clock generation problem, presenting the first PLL in n-type-only IGZO TFT technology and achieving bounded 1000 ppm frequency stability at $115\ \mu\text{W}$ across four published IGZO AMS systems.

Collectively, these contributions demonstrate that the co-design principle established in the preceding parts scales from individual circuit primitives all the way to deployable systems: each new layer of the stack—from reliable power supply through fault-tolerant ADC to optimized classifier and stable clock—builds directly on the techniques developed in the earlier parts, and each layer’s efficiency gains compound multiplicatively at the system level.

10. Cross-Layer Reliability for Flexible Healthcare Electronics

10.1. Motivation

Flexible healthcare systems must operate reliably in demanding environments—on the skin surface for wearables, or within the body for implantable devices—where exposure to mechanical stress, biological fluids, and temperature fluctuations accelerates device degradation beyond what is seen in laboratory conditions [142, 22]. Ensuring reliable operation requires a cross-layer approach that addresses three orthogonal challenges simultaneously: (i) stable power delivery despite load and process variation; (ii) fault-tolerant signal acquisition from the ADC interface; and (iii) passive, chipless sensing strategies that eliminate on-body active power for implantable applications [120].

The cross-layer nature of these challenges is important to emphasize. Supply voltage variation does not merely increase power consumption—a 100 mV shift in V_{DD} in IGZO circuits translates directly into comparator threshold shifts exceeding 1 LSB, so an unreliable power rail corrupts ADC conversion accuracy even in the complete absence of structural faults in the ADC itself. Conversely, structural faults in the Flash ADC propagate through the thermometer-code-to-binary encoder to produce incorrect digital outputs that the downstream classifier cannot distinguish from correct ones. And for implantable systems, both of these active-electronics challenges are compounded by the impossibility of integrating batteries or accessible test pins, demanding a completely passive sensing modality. This chapter presents the cross-layer reliability framework of [120], which combines a gain-boosted IGZO-TFT low-dropout voltage regulator (*ReFlex-LDO*) with integrated structural BIST, a hierarchical fault modeling and hardening methodology for IGZO-based Flash ADCs, and a passive RF-based sensing mechanism for implantable hydrocephalus shunt monitoring.

10.2. Reliable Power Delivery: ReFlex-LDO

10.2.1. Architecture

The *ReFlex-LDO* is an all-NMOS low-dropout regulator designed in the PragmatIC FlexIC Gen3 process, addressing the challenges of reduced intrinsic gain, high process variation, and the absence of PMOS devices in FE [120]. The complete schematic and layout are shown in Figure 10.1. The design retains the key stages of a classical LDO [132, 66] but adapts each for IGZO TFTs through four main stages.




Figure 10.1.: Schematic and post-layout view of the proposed gain-booster ReFlex-LDO [120]. The design includes a differential error amplifier (M5–M6), a second gain stage (M8–M9) that compensates for the limited intrinsic gain of IGZO TFTs, the pass transistor M_{pass} , a resistive feedback divider, and an on-chip output capacitor $C_0 = 0.95$ pF for dominant-pole compensation.

Error amplifier. The loop begins with the differential input pair (M5–M6), which senses the difference between the reference voltage V_{REF} and the divided feedback voltage V_{FB} . In silicon LDOs, this stage typically uses a PMOS differential pair for common-mode input range; here, an NMOS pair is used with a carefully chosen bias point to accommodate the limited V_{GS} range of IGZO TFTs.

Second gain stage. The error signal is fed into a second stage (M8–M9), which boosts both voltage gain and output swing to drive the large pass transistor M_{pass} efficiently. This gain-boosting stage is the critical architectural innovation: IGZO TFTs exhibit low intrinsic gain ($g_m r_o \approx 10$ –20) compared to silicon ($g_m r_o > 1000$), which would result in inadequate loop gain and poor line regulation if only a single-stage amplifier were used. The second stage multiplies the loop gain by its own voltage gain, restoring adequate regulation even with a pF-scale on-chip capacitor.

Pass device and feedback network. The pass transistor M_{pass} provides the load current and regulates V_{out} . A resistive feedback divider ($R_{\text{FB1}}-R_{\text{FB2}}$) samples V_{out} and generates V_{FB} , closing the regulation loop.

Stability and compensation. The design employs an output-capacitor-dominant-pole compensation strategy using $C_0 = 0.95 \text{ pF}$ at V_{out} . This places the dominant pole at the output node, ensuring a phase margin of over 45° across PVT corners without relying on ESR-zero compensation. The loop gain and phase response after post-layout extraction are shown in Figure 10.2(a), confirming stable operation.

10.2.2.DC and AC Performance

Figure 10.2 presents the DC regulation and AC characteristics of the ReFlex-LDO after post-layout RC extraction. Line regulation at the output stage shows V_{out} stable across all PVT corners over the full V_{DD} operating range; the internal reference V_{ref} is similarly stable, with less than 50 mV variation across all evaluated process corners and a Monte Carlo spread of 5000 samples within $\pm 3\sigma$ of the nominal value. Load regulation shows V_{out} remaining within specification across the full I_{out} range. The PSRR reaches 34 dB at 10 kHz, sufficient to attenuate supply ripple arising from switching activity in co-located digital logic. Transient response under supply step variations settles within 30 μs (fast step) and 120 μs (slow step), and the quiescent current is 31 μA , making the LDO suitable for battery-powered wearable patches operating under tight energy budgets.




Figure 10.2.: Electrical performance of the ReFlex-LDO after post-layout extraction [120]. (a) V_{out} versus V_{DD} across PVT corners, showing line regulation. (b-c) V_{ref} versus temperature including 5000 Monte Carlo simulations. (d) Loop gain and phase response versus frequency after post-layout extraction. (e) PSRR versus frequency. (f) V_{out} response under supply ripple at V_{DD} . (g) Line transient response under varying V_{DD} conditions.

10.2.3. Structural Delay-Based BIST

To detect in-field degradation without external test equipment, the ReFlex-LDO integrates a structural delay-based BIST following the ON/OFF-keying perturbation methodology of [79, 24]. The BIST principle is as follows: a controlled perturbation is injected at a designated node by momentarily bypassing a passive or active component using a pass transistor; the resulting change in V_{out} is observed as a digitally measurable time delay t_d via a chain of inverters followed by a counter. Faults that shift the DC operating point of the LDO alter t_d in a detectable way, allowing pass/fail decisions without any analog-to-digital conversion.

The fault detection decision rule uses a two-sided guardband computed from fault-free Monte Carlo characterisation:

$$\mathcal{G} = [\mu_{t_d} - 3\sigma_{t_d} - T_{\text{clk}}, \mu_{t_d} + 3\sigma_{t_d} + T_{\text{clk}}], \quad (10.1)$$

where μ_{t_d} and σ_{t_d} are the mean and standard deviation of the fault-free delay distribution computed from 5000 Monte Carlo samples at each PVT corner and bending radius, and T_{clk} is the counter clock period. A fault is detected when t_d falls outside $\mathcal{G}$; the production guardband is formed by the intersection of guardbands across all characterized corners to avoid false positives.

Two perturbation methods are used: bypassing passive components (resistors or capacitors) with a pass transistor, and bypassing active components by pulling a bias device's gate, which perturbs multiple stages simultaneously. Each injection point adds only one pass transistor and a digital enable signal, so the hardware overhead is proportional to the number of injection points. Eight strategically placed injection points are characterized: (1) feedback-leg bypass (R_0/R_{10}), (2) error-amplifier tail starve (M7), (3) pass-gate discharge (M_{pass}), (4) V_{ref} pull (R_8 shunt), (5) driver-node clamp (M6–M9), (6) load step at V_{out} , (7) shared bias-rail pull, and (8) V_{DD} micro-droop.

10.2.4. Fault Coverage and Overhead

Table 10.1 summarizes the fault coverage of the ReFlex-LDO BIST.

Table 10.1.: Fault coverage of the ReFlex-LDO under structural delay-based BIST [120]. Open-circuit: 1 M Ω –1 G Ω ; short-circuit: 1–100 Ω .

Block	Simulated faults			Detected faults			Coverage
	Total	Shorts	Opens	Total	Shorts	Opens	
LDO core	73	36	37	60	28	32	82.19%
BIST circuitry	62	30	32	57	28	29	91.94%
Parametric	5	—	—	4	—	—	80.00%
Total	140	66	69	121	56	61	86.42%

Coverage scales favorably with the number of injection points. With a single injection node, coverage is 86.4%; adding a second node raises it to 89.3%; three nodes reach 91.4%; and saturation occurs at seven to eight nodes with 95.7%, as shown in Table 10.2. The diminishing returns beyond five injection points indicate that most structural faults are observable through a limited set of strategically placed locations. At the optimal eight-node configuration, area overhead is only 4.24% and test power is 1.84 ms total at 5.56% power overhead, with a per-point test time of 12 μ s.

Table 10.2.: Coverage, area, and power overhead versus number of BIST injection points [120].

Injection points	1	2	3	4	5	6	7	8
Coverage (%)	86.4	89.3	91.4	92.9	94.3	95.0	95.7	95.7
Δ Area (%)	1.56	2.54	3.25	3.96	4.11	4.23	4.23	4.24
ΔP_{test} (%)	1.42	1.45	2.73	3.92	4.85	5.50	5.53	5.56
ΔT_{tot} (ms)	1.11	1.12	1.24	1.36	1.48	1.60	1.72	1.84

Bending effects. Coverage is evaluated across four mechanical conditions using an empirical bending model that perturbs TFT parameters as a function of bending radius: $\mu \leftarrow \kappa_{\mu}(r)\mu_0$, $R_{\text{sw}} \leftarrow \kappa_R(r)R_{\text{sw},0}$, and $V_{th} \leftarrow V_{th,0} + \Delta V_{th}(r)$ [120]. At flat condition, coverage is 95.5%; at 20 mm radius it decreases to 89.8%; at 10 mm radius it remains 92.4%; and at the extreme 5 mm radius it drops to 86.7%. The non-monotonic relationship between radius and coverage (90% at 20 mm but 92% at 10 mm) reflects the non-linear bending model: at certain radii, the induced V_{th} shift moves the LDO’s operating point into a region where fault-induced delay changes are more observable against the guardband.

Temperature effects. Coverage remains above 88% at both -20°C and 0°C , but decreases towards 82% at 80°C due to the negative temperature coefficient of the flexible

resistors, which shifts the delay distribution toward the guardband boundary. The combined bending-and-temperature effect at 85 °C, 5 mm radius represents the most adverse condition and motivates future work on temperature-adaptive guardband calibration.

10.3. Fault-Tolerant Flash ADC

10.3.1. Architecture and Simulation Setup

The Flash ADC targeted in this chapter is a 3-bit design comprising seven comparators (COM0–COM6) in a parallel array connected to a resistor ladder, followed by a thermometer-to-binary priority encoder, as shown in Figure 10.3. The architecture and the IGZO-TFT comparator schematic are illustrated in the same figure.




Figure 10.3.: 3-bit Flash ADC and IGZO-TFT comparator used in the cross-layer reliability framework [120]. (a) Flash ADC architecture: seven comparators in parallel with a resistor ladder, followed by a 7-to-3 priority encoder. (b) IGZO-TFT comparator schematic: 7 transistors implementing the differential comparison and output latch.

All simulations are conducted in Cadence Spectre with the PragmatIC Gen2 FlexIC PDK. $V_{DD} = 1$ V; the input signal is a 1 V sinusoid at 5 Hz. The baseline comparator transistor sizing is given in Table 10.3: the differential input pair (M0, M1) and output stage (M5, M6) use $W = 2 \mu\text{m}$, $L = 600$ nm, while the intermediate latch transistors (M3, M4) use $W = 5 \mu\text{m}$ to provide the higher drive strength needed in the middle stages. The resistor in each comparator is 255 M Ω ; the resistor ladder uses 1.7 M Ω per segment.

Table 10.3.: Baseline Flash ADC comparator and resistor ladder sizing [120].

Component	Element	Size
COM0–COM6	M0, M1, M2, M5, M6	$W = 2\ \mu\text{m}, L = 600\ \text{nm}$
COM0–COM6	M3, M4	$W = 5\ \mu\text{m}, L = 600\ \text{nm}$
COM0–COM6	R	255 M Ω
Resistor ladder	per segment	1.7 M Ω

10.3.2. Hierarchical Fault Injection Framework

The fault injection methodology, illustrated in Figure 10.4, follows an end-to-end hierarchical flow: single-fault characterisation builds the fault library; system-level propagation classifies each fault; multi-fault scenarios are then constructed from the classified fault pairs; and the identified vulnerability profile guides targeted design hardening.




Figure 10.4.: End-to-end hierarchical fault injection flow for IGZO-based Flash ADCs [120]. Single-fault injection and DNL classification build the vulnerability profile; multi-fault scenarios test design robustness; sensitive transistors guide targeted comparator hardening.

Structural defects are modeled as resistive opens (500 M Ω) and shorts (10 Ω) at each transistor terminal (gate, drain, source), yielding five fault types per transistor: OC_D, OC_G, OC_S, SC_DG, and SC_GS. The ADC output is evaluated via Differential Nonlinearity (DNL) across the full input range; faults are classified as *benign* ($\text{DNL} < 0.5\ \text{LSB}$), *marginal* ($0.5 \leq \text{DNL} < 1\ \text{LSB}$), or *catastrophic* ($\text{DNL} \geq 1\ \text{LSB}$). An automated Python framework coordinates netlist parsing, fault injection, parallel Spectre execution, and metric extraction; over 220 fault configurations are evaluated per design variant.

10.3.3. Fault Sensitivity Analysis

Table 10.4 presents the fault sensitivity heat map for the baseline Flash ADC under single-fault injection. Color coding: green = benign (DNL < 0.5), yellow = marginal (0.5–1.0), red = catastrophic (> 1.0).

Table 10.4.: Fault sensitivity heat map for the 3-bit Flash ADC under single-fault injection. Values are RMS DNL deviations. Green: benign (<0.5); yellow: marginal (0.5–1.0); red: catastrophic (>1.0). Data from [120].

Block	T-id	OC_D	OC_G	OC_S	SC_DG	SC_GS
ComFlash0	M3	0.30	0.29	0.28	1.82	0.62
ComFlash0	M4	0.26	0.00	1.32	0.29	0.13
ComFlash0	M5	0.28	2.58	0.24	1.02	0.00
ComFlash0	M6	0.28	0.24	0.31	1.02	0.25
ComFlash0	M2	0.27	0.66	0.22	0.00	0.61
ComFlash1	M3	0.29	0.42	0.61	1.46	0.56
ComFlash1	M5	0.39	2.06	0.35	0.63	0.59
ComFlash1	others	≤ 0.65 (marginal or benign)				
ComFlash2	M3	0.41	0.6	0.37	1.11	0.49
ComFlash2	M5	0.30	1.61	0.33	0.25	0.61
ComFlash2	others	≤ 0.73 (marginal or benign)				
ComFlash3	all	≤ 0.87 (marginal or benign)				
ComFlash4-6	M3	OC_G catastrophic, rest ≤ 0.84 (marginal or benign)				
ComFlash4-6	others	≤ 0.64 (Marginal or benign)				

The heat map reveals an important pattern: catastrophic failures (> 1.0) are concentrated in the first two comparators (ComFlash0 and ComFlash1), and specifically at M3 (OC_D, SC_DG), M4 (OC_S), M5 (OC_G, SC_DG), and M6 (SC_DG). The physical interpretation is as follows. M3 and M4 are the intermediate latch transistors; a short-circuit between their drain and gate (SC_DG) creates a diode-connected load that shifts the comparator's switching threshold by a full supply swing, producing a DNL error that misaligns one comparator output relative to the thermometer code by one or more levels. M5 (OC_G) disconnects the gate of the output-stage transistor, causing the output to float to a state independent of V_{in} —an intrinsically catastrophic failure. Comparators COM4–COM6 show only one catastrophic fault in the OC_G node of M3, while the rest of the codes are benign or marginal, outperforming all comparators is COM3, which shows no catastrophic outputs; its position in the thermometer code means that an error shifts one code boundary by a fractional LSB rather than causing a code inversion, producing only marginal or benign DNL. Overall, approximately 95% of single faults produce DNL ≤ 1 LSB (61.6%

benign, 33.5% marginal, 4.9% catastrophic), confirming that the baseline design is already relatively robust but has specific critical transistors that require targeted hardening.

10.3.4. Fault-Tolerant Comparator Variants

Based on the vulnerability profile, five comparator variants are proposed, each introducing one or two redundant transistors at the critical nodes identified above. The topologies of variants D1–D5 are illustrated in Figure 10.5.




Figure 10.5.: Proposed fault-tolerant comparator variants for the Flash ADC [120]. (a) D1: single redundant transistor M7 at the M3/M4 branch. (b) D2: single redundant transistor M8 at a complementary position. (c) D3: two redundant transistors (M7, M8) at the output latch (M5/M6). (d) D4: combined redundancy of D1, D2, and D3. (e) D5: a distinct configuration with two-level redundancy (M7–M12) providing the highest coverage for both OC and SC fault types.

The motivation for each variant follows directly from the sensitivity map. D1 and D2 each add a single transistor targeting the M3/M4 latch branch, which is responsible for most catastrophic SC_DG faults. D3 adds redundancy to the M5/M6 output stage, targeting the M5 OC_G catastrophic faults. D4 integrates all three redundancies but at the cost of increased power (the M5/M6 stage now has more parallel current paths, increasing static dissipation). D5 is redesigned from the ground up with two levels of redundancy (M7–M10 and M11–M12) that simultaneously protect against both OC and SC failure modes

in a more balanced topology; this is why D5 is the only variant that achieves significant improvement in SC fault coverage without sacrificing OC coverage.

An important material-physics insight is relevant here: in FE, resistors dominate the chip area rather than transistors, since each comparator's pull-up resistor (255 M Ω) occupies far more layout area than any individual TFT. Consequently, adding one or two transistors increases comparator area by only 3%—far less than the 10–20% overhead that would be seen in silicon, where transistors and metal routing jointly determine area. The same area-insensitivity of transistors in FE is what makes targeted redundancy a practical and low-overhead strategy in this technology.

10.3.5. Single- and Multi-Fault Results

Table 10.5 summarizes the fault coverage, area, and power for all variants. Under single-fault injection, D5 achieves the highest coverage at 97.7% (B+M), compared to 95.1% for the baseline. The improvement is not uniform across fault types: OC-fault coverage reaches 99.3% in D5 (because the redundant transistors provide alternative current paths when one connection is broken), while SC-fault coverage reaches 95% (because shorts between nodes can simultaneously disable multiple transistors and are harder to mitigate by structural redundancy alone).

Table 10.5.: Fault coverage, area overhead, and power change for all Flash ADC comparator variants. B = benign (DNL < 0.5); M = marginal (0.5–1.0); C = catastrophic (>1.0). Area and power values are relative to the baseline. Data from [120].

Design	Single injection		Multi injection		Area (%)	Power (%)
	B/M/C (%)	B+M (%)	B/M/C (%)	B+M (%)		
Baseline	61.6/33.5/4.9	95.1	29.7/54.9/15.4	84.6	0	0
D0	61.6/31.4/6.5	93.0	31.9/54.9/13.2	86.8	0	–11.2
D1	61.6/35.5/2.9	97.1	38.5/50.6/11.0	89.0	+1.6	–8.7
D2	71.0/25.7/3.3	96.7	29.7/57.1/13.2	86.8	+1.6	–8.6
D3	38.8/49.8/11.4	88.6	22.0/63.7/14.3	85.7	+3.2	–3.3
D4	64.1/31.8/4.1	95.9	39.6/46.2/14.3	85.7	+3.2	+7.4
D5	69.8/27.8/2.4	97.7	41.8/49.4/8.8	91.2	+3.1	–8.0

The multi-fault simulation set (100 double-fault scenarios per design) reveals the practical impact of the redundancy strategies through concrete examples. When combining a benign fault (COM2 M3 OC_D) with a marginal fault (COM6 M3 OC_D), the baseline produces marginal DNL; D1, D4, and D5 reduce this to benign. When combining two

marginal faults—COM3 M3 SC_DG and COM2 M3 OC_D—the baseline escalates to catastrophic DNL; D0, D1, D4, and D5 contain this at marginal. The most discriminating scenario pairs two marginal faults, COM0 M3 SC_GS and COM6 M3 OC_S: the baseline and all variants except D5 escalate to catastrophic; D5 alone maintains marginal, confirming that its two-level redundancy uniquely captures this cross-fault amplification path. In the most severe test pairing catastrophic and marginal faults (COM0 M3 SC_GS with COM6 M3 OC_G), only D5 limits the outcome to marginal rather than catastrophic, which is the defining advantage of its distinct topology. The physical layouts of the baseline comparator, D5 comparator, and corresponding full Flash ADC designs are shown in Figure 10.6; total area increases from 0.484 mm² (baseline) to 0.5 mm² (D5).




Figure 10.6.: Physical layout comparison of the baseline and D5 fault-tolerant Flash ADC [120]. (a) Baseline comparator (left) and full ADC (right): total area 0.484 mm². (b) D5 comparator (left) and full ADC (right): total area 0.500 mm² (+3.1% area overhead, –8% power). The resistor ladder and comparators dominate the footprint; added transistors contribute marginally to total area.

10.4. Passive Implantable Sensor

10.4.1. Clinical Motivation: Hydrocephalus Shunt Monitoring

Implantable electronic systems face a fundamental constraint: active on-body power is incompatible with long-term implantation due to risks from battery degradation, electrolyte leakage, and the surgical procedure required for replacement. The clinical target application is the monitoring of CSF flow in hydrocephalus shunt systems. Hydrocephalus is a neurological disorder requiring lifelong management via shunts that drain excess cerebrospinal fluid from the brain ventricles [73, 131]. Up to 40% of shunt systems fail within

two years of implantation [35], yet definitive diagnosis currently requires surgical exploration because imaging modalities (CT, MRI, cranial ultrasound) have limited reliability in detecting shunt malfunction. This motivates a fully passive, chipless monitoring mechanism that can assess shunt patency continuously without any on-board power or active electronics within the implant.

10.4.2. RF-Based Sensing Principle

The proposed system, shown in Figure 10.7, exploits the electrical conductivity of CSF itself as the sensing element. A monopole antenna and its associated bias line are fully integrated into the silicone catheter wall, with maximum width 1.6 mm to ensure compatibility within standard shunt catheters (inner diameter ≈ 1 mm, outer diameter ≈ 2 mm). A bias-line pad at one end of the antenna and an end-line pad at the other are positioned such that CSF flow through the shunt valve electrically connects them.




Figure 10.7.: Passive RF-based monitoring system for hydrocephalus shunt patency [120]. (a) System overview: external wearable transceiver interrogates the passively implanted antenna. (b) Valve closed (supine): sensing pads disconnected; antenna resonates at 915 MHz. (c) Valve open (upright): CSF electrically connects the pads; resonance shifts to 405 MHz. Persistence of 915 MHz resonance during upright posture indicates shunt malfunction.

The sensing principle exploits posture-driven CSF flow. When the patient is supine, the shunt valve is closed: the sensing pads are electrically disconnected and the antenna resonates at 915 MHz (ISM band), corresponding to the nominal idle state. When the patient is upright, gravitational pressure drives CSF flow through the open valve; the intrinsic electrical conductivity of CSF ($\sigma_{\text{CSF}} \approx 2 \text{ S/m}$) creates a conductive bridge between the pads, shifting the antenna resonance to 405 MHz (MedRadio band). An external wearable RF transceiver monitors this frequency shift as a change in backscattered power ΔP_{out} , enabling shunt patency assessment without any on-board power. Persistence of 915 MHz

resonance during upright posture indicates absence of CSF flow—a direct indication of shunt malfunction.

10.4.3. Prototype Fabrication and Experimental Validation

Prototype antennas were designed in ADS and fabricated on RO4003C substrate using an LPKF S104 milling system. The implanted antenna occupies $L = 51$ mm, $W = 1.58$ mm; the transceiver antenna is a broadband design ($L = 237.5$ mm, $W = 93.8$ mm) capable of operating at both frequencies simultaneously. Benchtop characterisation uses 0.9% saline solution (conductivity matched to CSF) to emulate open and closed valve states.

Saline impedance characterisation across ten pad-pair geometries (5 mm and 2.5 mm spacing, varying pad lengths 1–5 mm) confirms frequency-dependent reactive behaviour: at 405 MHz the average impedance is $1.277 + j1.459 \Omega$ (inductive), while at 915 MHz it is $1.929 - j1.515 \Omega$ (capacitive). Critically, increasing the saline volume from 10 μ L to 20 μ L produces minimal change across all geometries, demonstrating that the frequency response is governed by the conductive bridge geometry rather than the fluid volume—an important robustness property for real-world deployment where CSF flow rate varies.

The measured S_{11} return losses of the fabricated transceiver antenna are -19.3 dB at 405 MHz and -22.8 dB at 915 MHz, confirming adequate impedance matching across both operating bands. The simulated and measured ΔP_{out} between the two valve states is approximately 9.1 dB (valve closed) and 7.6 dB (valve open), both sufficient for reliable digital discrimination by the wearable transceiver. The use of two standardised frequency bands (ISM at 915 MHz and MedRadio at 405 MHz) provides an inherent calibration reference: the ratio between the two backscatter levels encodes valve state in a form that is insensitive to absolute link budget variations from patient anatomy or transceiver distance.

10.5. System Integration and Summary of Results

The three contributions described above form a coherent cross-layer reliability framework, illustrated conceptually as follows. The ReFlex-LDO provides a stable power rail (V_{out} regulated against process, temperature, and bending variation) for the wearable electronics module. The fault-tolerant Flash ADC (D5 variant) converts physiological signals with 97.7% single-fault and 91.2% multi-fault coverage, ensuring that structural defects in the comparator array do not propagate unchecked into the digital processing pipeline. The passive implantable sensor completes the system by monitoring shunt patency from the implant side without any active electronics within the body, interrogated by the wearable transceiver unit.

LDO coverage summary. Of 140 injected faults, 121 are detected (86.42% overall). The LDO core reaches 82.19% and the BIST circuitry achieves 91.94%. At flat condition with 5 or more injection points, coverage reaches 95.5%. Coverage degrades to 86.7% at 5 mm bending radius and 82% at 80 °C—conditions that represent the absolute extremes of wearable deployment and are in practice addressed by system-level scheduling of BIST cycles during rest periods.

Flash ADC coverage summary. The D5 variant achieves 97.7% single-fault and 91.2% multi-fault benign-plus-marginal coverage, compared to 95.1% and 84.6% for the baseline. OC-fault coverage reaches 99.3% and SC-fault coverage reaches 95%. The total Flash ADC area increases from 0.484 mm² to 0.500 mm² (+3.1%), and power decreases by 8.0% due to the resizing of M3/M4 in D5. These results confirm that targeted transistor-level redundancy—guided by a fault sensitivity heat map—is significantly more hardware-efficient than uniform redundancy: applying full triple-modular redundancy to all comparators would triple the comparator area, while D5 achieves comparable coverage at 3.1% overhead.

Passive sensor validation. Benchtop measurements confirm distinct frequency responses for the two valve states, with S_{11} differences exceeding 7.6 dB in received power between valve open and valve closed conditions at both operating frequencies. The response is volume-independent (10 μ L vs. 20 μ L saline yield identical results), confirming robustness against CSF flow rate variation.

10.6. Summary and Contributions

1. A gain-booster IGZO-TFT LDO (*ReFlex-LDO*) with integrated structural delay-based BIST, achieving up to 95.5% fault coverage at the nominal flat condition with 4.2% area overhead and 1.8 mW during test. The BIST uses a two-sided guardband decision rule calibrated from 5000 Monte Carlo samples, enabling reliable fault detection across temperature (−20 °C to 80 °C) and mechanical deformation (5–20 mm bending radius).
2. A hierarchical fault injection and mitigation framework for IGZO-based 3-bit Flash ADCs, combining an automated Python simulation pipeline with a fault sensitivity heat map. The analysis identifies M3 (SC_DG), M4 (OC_S), M5 (OC_G), and M6 (SC_DG) as the primary catastrophic fault locations. Five targeted comparator variants (D1–D5) are proposed; the D5 variant achieves 97.7% single-fault and 91.2% multi-fault benign-plus-marginal coverage (OC: 99.3%, SC: 95%), with only 3.1% area increase and 8% power reduction, representing a substantially more efficient solution than uniform TMR.
3. A fully passive, chipless RF sensing mechanism for hydrocephalus shunt monitoring, exploiting CSF conductivity to shift antenna resonance between 915 MHz (valve closed) and 405 MHz (valve open). Prototype validation confirms $\Delta P_{\text{out}} \geq$

7.6 dB between valve states and volume-independent response, enabling reliable shunt-patency assessment without any on-board power.

4. A unified cross-layer reliability framework that pairs passive implantable sensing with fault-tolerant wearable electronics, demonstrating that reliability in flexible healthcare systems requires co-design across power management, signal acquisition, and system-level monitoring layers simultaneously.

11. Design Space Exploration of Low-Power Wearable Stress Monitoring Classifiers

11.1. Motivation

The literature on wearable stress monitoring has focused predominantly on algorithmic solutions—identifying which ML model achieves the highest accuracy on a given physiological dataset—with limited attention to the hardware implications of deploying those models on resource-constrained wearable devices [142, 72]. General-purpose silicon microcontrollers, the dominant hardware platform for existing wearables, are rigid, expensive, and consume continuous power for both computation and communication, making them unsuitable for the disposable, conformable, low-power devices envisioned for FE healthcare applications.

Within FE, prior work had addressed flexible biosensors and ADCs separately, but the ML stress classifier—the critical component that must make real-time predictions from the digitized physiological signals—had not been systematically designed and characterized in flexible technology [8, 58]. This chapter presents the first comprehensive design space exploration (DSE) of low-power flexible stress monitoring classifiers, evaluating over 1200 classifier designs across three ML algorithms, three statistical feature selection methods, and four arithmetic precision levels, with unstructured pruning applied to MLPs.

11.2. System Architecture

The targeted flexible classification system is illustrated in Figure 11.1. Flexible biosensors capture physiological stress-related signals—primarily electrodermal activity (EDA), blood volume pulse (BVP), and skin temperature (TEMP)—which are digitized by per-sensor ADCs, processed by feature extractors to derive statistical features, and classified by an ML model that outputs a stress/no-stress prediction. The present work focuses specifically on optimizing the ML classifier within this pipeline.




Figure 11.1.: Overview of the flexible real-time stress monitoring classification system. Biosensors capture physiological signals; per-sensor ADCs digitize them; feature extractors compute statistical features; the ML classifier predicts stress level.

All classifiers are implemented as fully-parallel bespoke circuits [119], where trained model coefficients are hardwired into combinational logic, eliminating the need for memory and registers—which are costly in FE—and enabling the lowest possible area and power for each specific trained model.

11.3. DSE Framework

The proposed DSE framework systematically evaluates all combinations of the following design dimensions:

- **ML algorithms:** Decision Trees (DTs), Support Vector Machines (SVMs) with linear kernels, and Multi-Layer Perceptrons (MLPs) with ReLU activation.
- **Feature selection:** Three statistical algorithms—Double Input Symmetrical Relevance (DISR), Fisher Score, and Joint Mutual Information (JMI)—applied with feature counts ranging from 5 to 30.
- **Arithmetic precision:** Fixed-point representations at 4, 6, 8, and 10 bits, applied to all model coefficients.
- **MLP pruning:** Unstructured L2-norm pruning-aware retraining at sparsity levels of 20%, 50%, and 90%.

The algorithmic flowchart of the DSE is shown in Figure 11.2.




Figure 11.2.: Algorithmic flowchart of the proposed design space exploration. Feature selection, classifier training, pruning, and quantization compose a pipeline; hardware evaluation maps each design to its power consumption; Pareto-optimal accuracy–power solutions are extracted.

A custom standard cell library, optimized for ultra-low-power operation at 1 V using the PragmatIC FlexIC PDK (Gen2 Helvellyn 2.1.0), is developed and characterized as part of this work. Operating at 1 V—significantly below the commercial 3 V standard—reduces cell size and power while maintaining a $1\ \mu\text{s}$ maximum delay across all cells. The resistor-NMOS layouts of the 2-input NAND and NOR gates are illustrated in Figure 11.3.




Figure 11.3.: Resistor-NMOS layout of 2-input NAND gate optimized for 1 V operation, showing the pull-up resistor and n-type transistor arrangement.

11.4. Datasets

Two publicly available physiological stress datasets are evaluated: **WESAD**, which collects multimodal physiological signals from 17 participants using chest (RespiBAN) and wrist (Empatica E4) devices under controlled baseline and stress conditions for binary stress/no-stress classification [136]; and **AffectiveROAD**, which captures driving-scenario stress from 9 experienced drivers across 13 sessions using the Empatica E4 [65].

11.5. Results

Feature selection. No single feature selection algorithm dominates across all classifiers and datasets. For WESAD, DISR and Fisher contribute 67% and 71% of Pareto-optimal points for SVMs and MLPs, respectively.

DSE evaluation. Low-precision designs populate all Pareto fronts: for WESAD, 43% of Pareto-optimal classifiers use 4-bit precision, and only 18% require 8 bits or more. The accuracy–power evaluation across all classifiers and datasets is shown in Figure 11.4.




Figure 11.4.: Accuracy–power Pareto fronts from the full DSE across DT (a, d), SVM (b, e), and MLP (c, f) classifiers for the WESAD (a–c) and AffectiveROAD (d–f) datasets.

Most accurate designs. Decision Trees achieve up to 94% accuracy on WESAD with only $9\mu\text{W}$ power and 0.002 mm^2 area—within the range of flexible energy harvesters—while MLPs achieve 94% (WESAD) and 98% (AffectiveROAD) at higher hardware cost. Compared to silicon-based systems, flexible solutions achieve state-of-the-art accuracy with fully flexible, patch-based hardware and purely edge-based computation.

11.6. Summary and Contributions

1. The first comprehensive DSE of flexible ML classifiers for stress monitoring, evaluating over 1200 designs across three algorithms, three feature selection methods, four precision levels, and three pruning sparsities.
2. A custom 1 V standard cell library for the PragmatIC FlexIC PDK enabling ultra-low-power bespoke classifier synthesis.

3. Identification of Pareto-optimal flexible stress classifiers achieving state-of-the-art accuracy at as little as $9\text{ }\mu\text{W}$ and 0.002 mm^2 .

12. Feature-to-Classifier Co-Design for Mixed-Signal Smart Flexible Wearables for Healthcare at the Extreme Edge

12.1. Motivation

Even after applying the NSGA-II co-design of [2], feature extraction remains in the digital domain and accounts for 46% of total system area with ADCs responsible for a further 51% of power [139]. Digital implementations of even simple statistics—mean, standard deviation—require registers, adder trees, and shift logic consuming hundreds of thousands of μm^2 in IGZO FE [138, 135, 151, 157, 169]. This chapter addresses the feature extraction bottleneck by implementing feature extraction in the *analog domain*, before any digitization, using the first analog feature extractors in FlexIC technology. The key insight is that simple statistical functions (maximum, minimum, mean, sum) can be approximated continuously in analog circuitry with orders of magnitude lower area than their digital counterparts.

12.2. Mixed-Signal System Architecture

In the proposed system (Fig. 12.1), analog feature extractor circuits process biosensor signals continuously over a sliding window. At each window boundary, a multiplexer routes all computed analog features through a single shared SAR ADC; the digitized outputs are buffered and passed to a digital bespoke MLP classifier.




Figure 12.1.: Overview of the proposed flexible mixed-signal system architecture. Analog feature extractors process all sensor signals concurrently in the analog domain; a MUX routes their outputs through a shared SAR ADC; digitized features are buffered for the digital MLP classifier.

12.3. Analog Feature Extractor Circuits

Four analog circuits implement the core statistical features, all using n-type IGZO transistors with a custom two-stage op-amp. The circuits are illustrated in Figure 12.2.




Figure 12.2.: (a) Analog feature extractor circuits: peak-detector for Max/Min, op-amp integrator for Mean, scaling amplifier for Sum. (b) Two-stage IGZO op-amp with differential pair, common-source amplifier, level-shifting, and buffer stage.

Maximum. A peak-detector comprising a diode-connected n-type transistor and a hold capacitor charges to each new maximum; a windowing switch resets to V_l at each boundary [109].

Minimum. The identical topology with reversed diode orientation captures the minimum; reset pre-charges the capacitor to V_h .

Mean. An op-amp integrator continuously accumulates the input: $\mu \approx \frac{1}{R_M C_M} \int_0^T x(t) dt$, where T is the window duration. Reset returns the integrator to mid-scale V_{ms} .

Sum. Derived from the mean output via a non-inverting amplifier with gain $N = 1 + R_f/R_i$, where N is the number of samples.

Analog feature circuits achieve areas of 0.002 mm^2 (Max/Min) and $0.0084\text{--}0.0099 \text{ mm}^2$ (Mean/Sum) — **97% lower than digital equivalents.**

12.4. Flexible SAR ADC

A custom SAR ADC digitizes the multiplexed analog feature outputs. SAR is selected for moderate complexity, resolution scalability, and low static power [146, 45, 124, 8]. The n -bit SAR uses an R-2R ladder DAC, a sample-and-hold stage, and the op-amp comparator, implemented with the PragmatIC FlexIC Gen-3 digital cell library. A 4-bit ADC suffices for all evaluated datasets.

The SAR ADC schematic is shown in Figure 12.3.




Figure 12.3.: Schematic of the n -bit SAR ADC: sample-and-hold stage, n -bit SAR logic, R-2R ladder DAC, and gain-and-bias output stage.

12.5. Hardware-Aware Feature Selection via Differentiable Gating

The digital MLP is trained with a NAS-inspired stochastic feature gating layer embedded in the input stage. Each gate $z_i \in [0, 1]$ represents the inclusion probability of feature x_i , approximated by the Concrete distribution for gradient-based optimization [62, 168, 44, 106]. An area-informed regularization term adds analog feature hardware costs to the training loss:

$$\mathcal{L} = \mathcal{L}_{CE} + \lambda \sum_{i=1}^d \sigma(\log \alpha_i) \cdot c_i, \quad (12.1)$$

where c_i is the analog area cost of feature i and λ controls the accuracy–cost trade-off. After convergence, gates below threshold τ are removed; post-training Lottery-Ticket pruning [52] further reduces classifier area. The framework overview is shown in Figure 12.4.




Figure 12.4.: Algorithmic overview of the hardware-aware co-design framework. Stochastic gating embedded in training jointly minimizes analog feature extraction cost and maximizes accuracy; post-training Lottery-Ticket pruning further reduces digital classifier area.

12.6. Results

Analog circuit fidelity. SPICE outputs for all four analog features closely track the software reference values for the SPD dataset, with NMSE ≤ 0.005 . Simulation results are shown in Figure 12.5.




Figure 12.5.: SPICE analog output (blue) versus ideal software-computed feature values (orange) for the SPD accelerometer x-axis: (a) Max, (b) Min, (c) Mean, (d) Sum.

System-level results. Classification accuracy remains within 3% of ideal floating-point results despite analog feature extraction error. Compared to the most efficient purely-digital solution from [2], the framework achieves an average **48× area reduction** and **70× lower energy per inference**. Power stays below 20.3 mW and energy per inference below 1 μ J. The hardware-aware training simultaneously delivers approximately 8% higher classification accuracy than hardware-unaware statistical feature selection.

12.7. Summary and Contributions

1. The first analog feature extractors (Max, Min, Mean, Sum) in FlexIC IGZO technology, achieving 97% lower area than digital equivalents.
2. A holistic feature-to-classifier co-design framework combining analog feature extraction, a custom SAR ADC, NAS-inspired hardware-aware feature gating, and Lottery-Ticket pruning.
3. 48× area and 70× energy-per-inference reductions versus the best purely-digital co-design, with classification accuracy within 3% of floating-point software results.

13. Low-Power PLL-Based Clock Stabilization for Flexible IGZO AMS Systems

13.1. Motivation

Clock generation is commonly neglected in published IGZO AMS system designs: most reports exclude the clock generator from their power budget, implicitly assuming an off-chip source or a free-running ring oscillator (RO). This assumption has two critical consequences. First, when a state-of-the-art IoT CMOS PLL is integrated into representative IGZO systems, the clock alone consumes 29%–89% of total system power [91]—far exceeding the 25% clock-network share typical of silicon systems. Second, free-running ROs suffer from unbounded PVT drift: in the kilohertz operating range of FE AMS systems, moderate threshold-voltage shifts translate into frequency errors exceeding 30%, directly corrupting ADC sampling accuracy [120].

This chapter presents the first PLL designed for n-type-only IGZO TFT technology—and the first closed-loop clock stabilization solution in any metal-oxide TFT technology. The PLL provides bounded 1000 ppm frequency accuracy with 2.24 ns rms jitter at 115 μ W and 0.009 mm², replacing unbounded drift with reference-anchored stability at negligible system overhead.

13.2. Proposed PLL Architecture

The PLL adopts a minimal charge-pump topology adapted for n-type-only IGZO, comprising a Phase-Frequency Detector (PFD), Charge Pump (CP), Low-Pass Filter (LPF), and ring-oscillator VCO [57, 129]. The complete schematic is shown in Figure 13.1.




Figure 13.1.: Proposed PLL schematic. Left: PFD implemented from two pseudo-CMOS dynamic D flip-flops [38, 107]. Centre: charge pump and first-order RC low-pass filter. Right: three-stage ring-oscillator VCO with resistive pseudo-CMOS stages and V_{ctrl} -gated bias network.

Two design principles govern the architecture. **Frequency co-design:** the VCO is retargeted to any frequency from 1 kHz to 300 kHz by adjusting only the load capacitance C , with all transistor sizes and sub-blocks unchanged. **Overhead minimality:** the PLL must contribute negligible area and power relative to the host system, preserving feasibility of battery-powered, conformable patches [126, 152, 139].

13.2.1. Ring-Oscillator VCO

Three cascaded resistive pseudo-CMOS stages (pull-up resistors R_1 – $R_7 = 172.5 \text{ k}\Omega$, n-type TFT pull-downs) form the ring oscillator. A V_{ctrl} -gated bias network modulates stage current and thus oscillation frequency. Load capacitance C is the sole retargeting parameter, spanning 0.27 pF (300 kHz) to 52 pF ($\approx 1 \text{ kHz}$) in a near-linear log-log relationship.

13.2.2. PFD and Charge Pump

The PFD is implemented with pseudo-CMOS dynamic D flip-flops, avoiding p-type devices [38, 107]. A tri-state topology enables lock acquisition from the large ($> 30\%$) initial frequency offsets caused by IGZO PVT spread. The charge pump uses a zero-quiescent-bias design, eliminating static current paths and enabling sub-0.2 mW total power. Loop bandwidth is set to approximately $f_{ref}/10$, simultaneously suppressing TFT flicker noise, reducing charge-pump switching activity, and tracking slow PVT drift.

13.3. System Integration and Target Specifications

The PLL is validated against four published IGZO AMS systems spanning 1 kHz to 300 kHz (Table 13.1), illustrated conceptually in Figure 13.2.

Table 13.1.: Target IGZO AMS systems for PLL integration ($V_{DD} = 3$ V).

Reference	f_{clk}	Area (mm ²)	Power (mW)
FlexRV [126]	300 kHz	0.59–1.02	0.71–0.91
Co-processor [152]	150 kHz	2.41–2.44	1.49–1.53
F2C system [139]	10 kHz	0.20–27.6	0.14–20.3
Flexible SAR ADC [8]	≈ 1 kHz	9.4	2.93




Figure 13.2.: Conceptual integration of the proposed PLL within a flexible AMS platform. A single PLL topology provides clock signals to multiple processing blocks at distinct frequency ranges; only the VCO load capacitor is adjusted for each target.

13.4. Results

Steady-state performance. At the 300 kHz operating point (27 °C, TT corner, $V_{DD} = 3$ V), total power is 115.4 μ W (VCO: 110.6 μ W, PFD + CP/LPF: 4.8 μ W), active area is 0.00879 mm², rms jitter is 2.24 ns, and long-term frequency accuracy is 1000 ppm.

Frequency scalability. The VCO output frequency as a function of load capacitance is shown in Figure 13.3, confirming a monotonic, near-linear log-log relationship across

three decades. Retargeting from 300 kHz to 1 kHz requires only changing C from 0.27 pF to 52 pF.




Figure 13.3.: Open-loop VCO output frequency versus load capacitance C (log-log scale). Monotonic, near-linear relationship spanning 1 kHz to 300 kHz achieved by varying C from 0.27 pF to 52 pF.

PVT robustness. Across 27 PVT combinations (3 voltages $\times$ 3 corners $\times$ 3 temperatures), the PLL locks reliably at 27 °C and 45 °C in 8–9 of 9 corner combinations within $\pm 7.5\%$ of target. PVT sweep results are shown in Figure 13.4.





Figure 13.4.: Locked output frequency (top) and total power (bottom) across 27 PVT combinations: 3 temperatures $\times$ 3 process corners $\times$ 3 supply voltages. The PLL acquires lock reliably at 27°C and 45°C; at 0°C IGZO mobility is insufficient for 300 kHz operation.

System power impact. Fig. 13.5 compares the system power breakdown when each target is clocked by the proposed IGZO PLL versus this CMOS reference. The proposed IGZO PLL reduces clock power from 29%–89% (CMOS PLL) to 5%–52% of total system power across all four integration targets—up to 5.8 $\times$ lower. Compared to prior FE clocking solutions (free-running ROs and open-loop VCOs operating at 5–25 V), the proposed PLL reduces power by more than 400 $\times$ and footprint by more than 2000 $\times$, while being the only FE solution providing bounded frequency accuracy.



†Classifier at highest-accuracy operating point (0.14 mW).

Figure 13.5.: Absolute system power when each target is clocked by a SoA IoT CMOS PLL [91] (left bar) versus the proposed IGZO PLL (right bar). The system power (blue) is identical in both bars; only the clock contribution differs. Percentages above each bar indicate the clock's share of total power.

13.4.1. Comparison with Prior Clocking Approaches

Table 13.2 compares the proposed design against reported oscillator-based clocking solutions in FE. Prior works employ free-running ROs or open-loop VCOs operating at 5–25 V and consuming 0.4–50 mW—one to three orders of magnitude above the 0.115–0.153 mW of this design. None reports jitter or frequency accuracy, since these metrics are not meaningful for open-loop oscillators with inherently unbounded drift. Three distinctions stand out. First, this is the only FE clocking solution operating at 3 V, directly matching the supply of all target IGZO AMS systems and eliminating any level-shifting overhead. Second, it is the only solution providing bounded frequency accuracy (1000 ppm), replacing open-loop drift with reference-anchored stability—a property that is simply absent, not merely worse, in all prior FE works. Third, our solution presents the smallest footprint, with a reduction of more than 2000× w.r.t to VCO alternatives, and more than 400× w.r.t to RO solutions.

Table 13.2.: Comparison with prior clocking solutions in FE.

Work	Freq (kHz)	Power (mW)	Area (mm²)
[80] IGZO VCO	0.33–0.56	1.3	N/R
[81] SWCNT VCO	0.85–2.05	≈0.4	17.5
[34] IGZO RO	1820–6510	≈43	0.13
[61] IGZO RO	810	≈50	0.208
[143] IGZO RO	406–2100	13–43	1.6–4.5
This work IGZO PLL	1–300	0.115–0.153	0.008–0.03

13.5. Summary and Contributions

1. The first PLL in n-type-only IGZO TFT technology, achieving 1–300 kHz operation at 3 V with 0.115–0.153 mW and 0.008–0.028 mm².
2. A frequency-scalable architecture retargeted via VCO load capacitance alone, validated across four published IGZO AMS systems spanning three frequency decades.
3. More than 400× power and 2000× footprint reduction versus prior FE clocking solutions, with bounded 1000 ppm frequency accuracy replacing free-running oscillator drift.

14. Conclusions and Future Work

14.1. Summary of Contributions

This thesis has addressed the central challenge of enabling efficient, reliable, and application-ready analog and mixed-signal computing in Flexible Electronics. The four research directions pursued over three years collectively advance the state of the art along a coherent trajectory: from fundamental analog circuit primitives, through co-designed ADC–classifier systems, to fault-tolerant building blocks, complete healthcare wearable systems, and the clock generation infrastructure that underpins them. The unifying thread across all contributions is the hardware–software co-design principle: incorporating circuit-level models of error, cost, and variability into the training and optimization loop consistently outperforms approaches that treat hardware as a fixed post-training constraint.

Analog function approximation (Part I) established that nonlinear computation can be performed directly in the IGZO analog domain using a reusable library of five Analog Building Blocks. The $125\times$ area reduction over 8-bit digital splines confirmed that, for functions expressible as low-order polynomials, analog computation is substantially more efficient than digital computation in FE—a non-obvious result given the variability and limited precision of IGZO devices. The AKAN hardware–software co-design framework demonstrated that incorporating circuit-level error models into KAN training is not merely beneficial but necessary: ignoring hardware error leads to models that converge to coefficient regions associated with large hardware non-idealities, producing suboptimal implementations that could be substantially improved by closing the loop between simulation and training. Up to $5.5\times$ improvement in the area $\times$ power product, with average gains of approximately 30% across four benchmark datasets, validated this principle quantitatively.

ADC co-design (Part II) demonstrated that the ADC is not a fixed hardware constant but a primary co-design variable. Jointly optimizing ADC comparison levels with classifier decision boundaries—via NSGA-II for Binary Search ADC pruning, or via Artificial Bee Colony optimization for non-uniform threshold placement—consistently reduces ADC hardware cost by an order of magnitude without degrading classification accuracy. A notable finding across both contributions is that restricting the input representation through pruning or non-uniform quantization frequently *improves* accuracy rather than merely preserving it, because the reduced codebook implicitly regularizes the MLP against

noise in the continuous input domain. These results shift the design paradigm: ADC resolution and threshold placement should be learned jointly with classifier weights, not specified a priori.

Reliability and fault tolerance (Part III) addressed the elevated defect densities of IGZO manufacturing through targeted, hardware-aware strategies. The fault-sensitivity heat map for Binary Search ADCs revealed a structural asymmetry—early-stage comparators are catastrophically sensitive, whereas late-stage comparators exhibit 0% multi-fault coverage in the baseline—that would not have been identified by standard uniform redundancy. The selective redundancy approach that followed from this analysis recovers the majority of fault coverage (from 60% to 92% single-fault, from 34% to 77.3% multi-fault) at 4.2% area and 6% power overhead. The delay-based BIST demonstrated that fault coverage exceeding 90% can be achieved on-chip without external test equipment—a prerequisite for FE deployment—by exploiting inter-stage delay signatures of a low-power ring oscillator. The frequency-domain parametric fault detection methodology demonstrated that ratiometric RMS spectral signatures maintain fault separability across global PVT variation without any analog-to-digital conversion in the test path, directly addressing the test-invalidation problem inherent in conventional ADC-based monitoring approaches.

Healthcare applications and clock generation (Part IV) showed how the preceding contributions integrate into complete, deployable systems. The progression from a single-component DSE (classifier only) through full end-to-end co-design (ADC + feature extractor + classifier) and mixed-signal integration (analog feature extractors + SAR ADC + digital classifier) illustrates the cumulative power of the co-design principle when applied consistently across all system layers: each step in this progression yields a multiplicative, not additive, improvement in system efficiency. The PLL contribution addressed a blind spot in the FE literature: clock generation had been universally neglected, yet an unoptimized clock source can consume 90% of total system power. Providing the first PLL in IGZO technology—achieving 1000 ppm frequency accuracy at 115 μW in 0.009 mm², validated across four published IGZO AMS systems—establishes clock generation as a viable, low-overhead system component.

14.2. Open Problems and Future Research Directions

Several promising directions extend beyond the scope of this thesis.

Larger-scale analog KAN integration. The AKAN results demonstrate up to $5.5\times$ area-power improvement through two-level pruning, but the framework currently targets univariate functions. Extending AKANs to multivariate functions—where the Kolmogorov-Arnold representation theorem is most powerful—requires routing analog signals through multiple layers of splines, raising new challenges in signal range management, inter-layer

coupling, and systematic error accumulation across layers. The NMPE stability demonstrated across 27 PVT corners ($\Delta\text{NMPE} \leq 0.74$ pp at 27–40 °C) is an encouraging foundation for such an extension.

Online reconfigurability. All contributions in this thesis treat the hardware design as fixed post-fabrication. FE’s low non-recurring engineering cost makes it feasible to produce multiple variants of a design for different tasks; however, a single reconfigurable circuit that can adapt its ADC thresholds, non-uniform resistor ladder, or classifier weights in the field—without full redesign—would dramatically extend the practical utility of flexible systems. Emerging memristive components in FE, while not yet in the PDKs used in this thesis, could in principle provide online resistance adjustment.

Multi-sensor fusion and system-on-foil integration. The healthcare applications target individual physiological signals (primarily EDA and BVP). Clinical stress monitoring benefits from fusing multiple modalities—respiration, temperature, heart rate variability—simultaneously. Integrating multiple analog processing chains, Binary Search ADCs, and a shared classifier on a single FlexIC while managing inter-channel interference and synchronization across channels clocked by the proposed PLL is an open system integration challenge that the individual contributions of this thesis are now equipped to tackle jointly.

In-field reliability monitoring. The BIST architecture developed in Part III enables one-shot post-manufacturing testing and periodic in-field checks. A more ambitious goal is continuous, background reliability monitoring—using the frequency-domain Spectral signature as a real-time degradation indicator rather than a pass/fail comparator—which would enable predictive maintenance and adaptive fault mitigation in deployed wearable patches. The monotonic, PVT-robust spectral signatures demonstrated in Chapter 9 are directly suited for this extension.

Layout-aware training. Post-layout simulation under PVT conditions was incorporated for the Binary Search ADC (Chapter 6) and confirmed the importance of parasitic effects from metal routing and device geometry. Systematically incorporating post-layout RC extraction into the training loop of hardware-aware frameworks—so that models are optimized against post-layout, not schematic-level, hardware behavior—remains an open co-design challenge. The 6-minute average runtime of the ABC-based non-uniform quantization framework suggests that layout-in-the-loop optimization is feasible within practical design iteration cycles.

Adaptive thresholds for the frequency-domain Signature for Fault Detection. The fault coverage analysis of Chapter 9 revealed that fast-process corners exhibit paradoxically reduced relative sensitivity, because their higher absolute VCO gain compresses the normalized signature change for small faults. Per-corner adaptive thresholds—set automatically from the one-time post-manufacturing calibration measurement—would maximize detection capability across all corners without increasing false-alarm rates. This extension requires no new hardware and could be implemented entirely in the microcontroller already present in the wearable system.

Advancing towards system-level closure. The PLL, BIST, analog feature extractors, Binary Search ADC, and classifier contributions developed in this thesis have each been validated individually and, in several cases, against published partner systems. Integrating all components into a single validated FlexIC system-on-foil—sensor, analog front-end, Binary Search ADC, analog feature extractor, digital classifier, distributed BIST, and PLL clock—would constitute a significant system-level demonstration and validate the coherence of the co-design methodology across all layers simultaneously. The results presented in this thesis establish the individual building blocks, the co-design frameworks, and the validation methodology that such a demonstration would require.

Bibliography

- [1] F. Afentaki et al. “Exploration of low-power flexible stress monitoring classifiers for conformal wearables”. In: *Proceedings of the International Symposium on Low Power Electronics and Design (ISLPED)*. 2025.
- [2] F. Afentaki et al. *Real-Time Stress Monitoring with Flexible Electronics: A Co-Design Approach from ADCs to Feature Extractors and Classifiers*. 2026.
- [3] Florentia Afentaki et al. “Bespoke Approximation of Multiplication-Accumulation and Activation Targeting Printed Multilayer Perceptrons”. In: *Int. Conf. on Computer Aided Design (ICCAD)*. 2023, pp. 1–9. DOI: 10.1109/ICCAD57390.2023.10323613.
- [4] Florentia Afentaki et al. “Design and Optimization of Mixed-Kernel Mixed-Signal SVMs for Flexible Electronics”. In: *Design Automation and Test in Europe (DATE)*. 2026.
- [5] Florentia Afentaki et al. “Lightweight Fault Resilient Flexible Flash-ADCs”. In: *IEEE 32nd International Symposium on On-Line Testing and Robust System Design*. 2026.
- [6] Florentia Afentaki et al. “Reducing ADC Front-End Costs During Training of On-Sensor Printed Multilayer Perceptrons”. In: *IEEE Embedded Systems Letters* 16.4 (2024), pp. 353–356. DOI: 10.1109/LES.2024.3447412.
- [7] Arif U Alam et al. “Fruit Quality Monitoring with Smart Packaging”. In: *Sensors* 21.4 (2021), p. 1509.
- [8] Feras Alkhalil et al. “Flexible SAR ADC with resistive DAC for Conformable On-Body Sensing Applications”. In: *2022 IEEE Biomedical Circuits and Systems Conference (BioCAS)*. 2022, pp. 110–114.
- [9] K. Arabi, B. Kaminska, and M. Sawan. “On-chip testing of data converters using static parameters”. In: *IEEE Transactions on VLSI Systems* 6.3 (1998), pp. 413–424.

- [10] Jean Arlat et al. “Fault injection for dependability validation: a methodology and some applications”. In: vol. 16. 2. 1990, pp. 166–182. doi: 10.1109/32.44380.
- [11] Giorgos Armeniakos et al. “Co-Design of Approximate Multilayer Perceptron for Ultra-Resource Constrained Printed Circuits”. In: *IEEE Transactions on Computers* 72.9 (2023), pp. 2717–2725. doi: 10.1109/TC.2023.3251863.
- [12] Giorgos Armeniakos et al. “Model-to-Circuit Cross-Approximation For Printed Machine Learning Classifiers”. In: *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 42.11 (2023), pp. 3532–3544. doi: 10.1109/TCAD.2023.3258668.
- [13] Giorgos Armeniakos et al. “Hardware Approximate Techniques for Deep Neural Network Accelerators: A Survey”. In: *ACM Comput. Surv.* 55.4 (2022). doi: 10.1145/3527156.
- [14] Giorgos Armeniakos et al. “On-Sensor Printed Machine Learning Classification via Bespoke ADC and Decision Tree Co-Design”. In: *2024 Conference & Exhibition on Design, Automation & Test in Europe (DATE)*. 2024, pp. 1–6.
- [15] Alvaro Ayuso-Martinez et al. “Analog Implementation of a Spiking System for Performing Arithmetic Logic Operations on Mixed-Signal Neuromorphic Processors”. In: *Advanced Intelligent Systems* 6.2 (2024), p. 2400524. doi: 10.1002/aisy.202400524.
- [16] O. Bajdechi, G.E. Gielen, and J.H. Huijsing. “Systematic design exploration of delta-sigma ADCs”. In: *IEEE Transactions on Circuits and Systems I: Regular Papers* 51.1 (2004), pp. 86–95. doi: 10.1109/TCSI.2003.821293.
- [17] Ratul Baruah, Hocheon Yoo, and Eun Lee. “Interconnection Technologies for Flexible Electronics: Materials, Fabrications, and Applications”. In: *Micromachines* 14 (2023), p. 1131. doi: 10.3390/mi14061131.
- [18] Timoteo García Bertoa, Giulio Gambardella, et al. “Fault-Tolerant Neural Network Accelerators With Selective TMR”. In: *IEEE Design & Test* 40.2 (2023), pp. 67–74.

- [19] Shah Palash Manish Bhai and Dattatraya N. Gaonkar. “Design of binary search ADC using N comparators”. In: *2016 IEEE First International Conference on Control, Measurement and Instrumentation (CMI)*. 2016, pp. 499–502.
- [20] Gulafsha Bhatti et al. “Flexible Electronics: A Critical Review”. In: *Interconnect Technologies for Integrated Circuits and Flexible Electronics*. Ed. by Yash Agrawal, Kavicharan Mummaneni, and P. Uma Sathyakam. Singapore: Springer Nature Singapore, 2024, pp. 221–248.
- [21] John Biggs et al. “A natively flexible 32-bit Arm microprocessor”. In: *Nature* 595.7868 (2021), pp. 532–536.
- [22] Bülent Bilgehan, Ozlem Sabuncu, and Lara Kayed. “General probability distribution model for wireless body sensors in the medical monitoring system”. In: *Biomedical Signal Processing and Control* 77 (2022), p. 103777. DOI: 10.1016/j.bspc.2022.103777.
- [23] Bora Bilgic and Sule Ozev. “Low-cost structural monitoring of analog circuits for secure and reliable operation”. In: *IEEE Design & Test* 40.4 (2023), pp. 5–16.
- [24] Bora Bilgic and Sule Ozev. “Performance degradation monitoring for analog circuits using lightweight built-in components”. In: *2022 IEEE 40th VLSI Test Symposium (VTS)*. IEEE, 2022, pp. 1–7.
- [25] Swarup Biswas et al. “Emerging Energy Harvesters in Flexible Bioelectronics: From Wearable Devices to Biomedical Innovations”. In: *Small Science* 4 (2024).
- [26] J. Blank and K. Deb. “pymoo: Multi-Objective Optimization in Python”. In: *IEEE Access* 8 (2020), pp. 89497–89509.
- [27] Nathaniel Bleier et al. “Printed Microprocessors”. In: *Proceedings of the ACM/IEEE 47th Annual International Symposium on Computer Architecture*. ISCA '20. IEEE Press, 2020, pp. 213–226.
- [28] Yvan Bonnassieux et al. “The 2021 flexible and printed electronics roadmap”. In: *Flexible and Printed Electronics* 6.2 (2021), p. 023001.
- [29] Tech Briefs. *For Power-Efficient Wireless Sensing, Start with Analog Processing*. 2020.
- [30] Gabriel Cadilha Marques et al. “Digital power and performance analysis of inkjet printed ring oscillators based on electrolyte-gated oxide electronics”. In: *Applied Physics Letters* 111.10 (2017), p. 102103.

- [31] Indranil Chakraborty et al. “Resistive Crossbars as Approximate Hardware Building Blocks for Machine Learning: Opportunities and Challenges”. In: *Proceedings of the IEEE* 108.12 (2020), pp. 2276–2310. DOI: 10.1109/JPROC.2020.3003007.
- [32] Y. Chen and E. Keogh. *ECG5000: A 20-hour Long ECG Dataset for Classification*. 2008. URL: <https://www.timeseriesclassification.com/description.php?Dataset=ECG5000>.
- [33] Yen-Lung Chen et al. “Simultaneous Optimization of Analog Circuits With Reliability and Variability for Applications on Flexible Electronics”. In: *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 33.1 (2014), pp. 24–35. DOI: 10.1109/TCAD.2013.2282757.
- [34] Yuanfeng Chen et al. “High-Speed Pseudo-CMOS Circuits Using Bulk Accumulation a-IGZO TFTs”. In: *IEEE Electron Device Letters* 36.2 (2015), pp. 153–155. DOI: 10.1109/LED.2014.2379700.
- [35] Prashant Chittiboina et al. “Implant and patient characteristics associated with fatal shunt malfunction in patients with hydrocephalus”. In: *Journal of Neurosurgery: Pediatrics* 14.5 (2014), pp. 533–538. DOI: 10.3171/2014.7.PEDS1490.
- [36] Minwoo Choi et al. “Flexible active-matrix organic light-emitting diode display enabled by MoS₂ thin-film transistor”. In: *Science Advances* 4.4 (2018), eaas8721.
- [37] Claudionor N Coelho et al. “Ultra low-latency, low-area inference accelerators using heterogeneous deep quantization with QKeras and hls4ml”. In: *arXiv preprint arXiv:2006.10159* (2020), p. 108.
- [38] Júlio Costa et al. “Evolving into an Era of Natively Flexible Smart Systems”. In: *SID Symposium Digest of Technical Papers*. Pragmatic Semiconductor, 2023.
- [39] K. Deb et al. “A fast and elitist multiobjective genetic algorithm: NSGA-II”. In: *IEEE Transactions on Evolutionary Computation* 6.2 (2002), pp. 182–197. DOI: 10.1109/4235.996017.
- [40] Analog Devices. *Power Optimization Techniques for Low Power Signal Chain Applications*. 2022.

- [41] Alejandro Diaz-Sanchez et al. "A four quadrant high-speed CMOS analog multiplier based on the flipped voltage follower cell". In: *AEU - International Journal of Electronics and Communications* 130 (2021), p. 153582.
- [42] Dheeru Dua and Casey Graff. *UCI Machine Learning Repository*. 2017.
- [43] Fierce Electronics. *Analog Processing Improves Battery Life and Sensor Intelligence In the IoT*. 2020.
- [44] Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. "Neural Architecture Search: A Survey". In: *Journal of Machine Learning Research* 20.55 (2019), pp. 1–21.
- [45] Fabrizio Erario et al. "Design of an ultra-low power time interleaved SAR converter". In: 2008, pp. 245–248. DOI: 10 . 1109 / RME . 2008 . 4595771.
- [46] Osman Emir Erol et al. "On-chip measurement of bandgap reference voltage using a small form factor VCO based zoom-in ADC". In: *2015 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. 2015, pp. 1559–1562. DOI: 10.7873/DATE.2015.1003.
- [47] Osman Emir Erol and Sule Ozev. "Knowledge- and Simulation-Based Synthesis of Area-Efficient Passive Loop Filter Incremental Zoom-ADC for Built-In Self-Test Applications". In: *ACM Trans. Des. Autom. Electron. Syst.* 24.1 (2018).
- [48] EURORACTICE. *Flexible Electronics*. Accessed: 2025-02-28. 2025. URL: <https://europractice-ic.com/technologies/flexible-electronics/>.
- [49] Md Meftahul Ferdaus et al. *KANICE: Kolmogorov-Arnold Networks with Interactive Convolutional Elements*. 2024. arXiv: 2410 . 17172 [cs.CV].
- [50] Manuel Fernández-Delgado et al. "Do We Need Hundreds of Classifiers To Solve Real World Classification Problems?" In: *The Journal of Machine Learning Research* 15.1 (2014), pp. 3133–3181.
- [51] R. A. Fisher. *Iris*. UCI Machine Learning Repository. 1936. DOI: 10 . 24432/C56C76.
- [52] Jonathan Frankle and Michael Carlin. "The Lottery Ticket Hypothesis: Finding Sparse, Trainable Neural Networks". In: *International Conference on Learning Representations (ICLR)*. 2019.

- [53] E.G. Friedman. "Clock distribution networks in synchronous digital integrated circuits". In: *Proceedings of the IEEE* 89.5 (2001), pp. 665–692. DOI: 10.1109/5.929649.
- [54] Jiuwei Gao et al. "Ultra-Robust and Extensible Fibrous Mechanical Sensors for Wearable Smart Healthcare". In: *Advanced Materials* 34.20 (2022), p. 2107511.
- [55] Wei Gao et al. "Flexible Electronics toward Wearable Sensing". In: *Accounts of Chemical Research* 52 (2019). DOI: 10.1021/acs.accounts.8b00500.
- [56] Wei Gao et al. "Fully integrated wearable sensor arrays for multiplexed in situ perspiration analysis". In: *Nature* 529 (2016), pp. 509–514.
- [57] Floyd M. Gardner. "Charge-pump phase-lock loops". In: vol. 28. 11. 1980, pp. 1849–1858. DOI: 10.1109/TCOM.1980.1094619.
- [58] Carmine Garripoli et al. "Analogue Frontend Amplifiers for Bio-Potential Measurements Manufactured With a-IGZO TFTs on Flexible Substrate". In: *IEEE Journal on Emerging and Selected Topics in Circuits and Systems* 7.1 (2017), pp. 60–70. DOI: 10.1109/JETCAS.2016.2616723.
- [59] Carmine Garripoli et al. "15.3 An a-IGZO asynchronous delta-sigma modulator on foil achieving up to 43dB SNR and 40dB SNDR in 300Hz bandwidth". In: *2017 IEEE International Solid-State Circuits Conference (ISSCC)*. 2017, pp. 260–261. DOI: 10.1109/ISSCC.2017.7870360.
- [60] Enrico Genco et al. "An EMG Interface Comprising a Flexible a-IGZO Active Electrode Matrix and a 65-nm CMOS IC". In: *IEEE Journal of Solid-State Circuits* 58.11 (2023), pp. 3138–3149. DOI: 10.1109/JSSC.2023.3274709.
- [61] Di Geng, Dong Han Kang, and Jin Jang. "High-Performance Amorphous Indium–Gallium–Zinc–Oxide Thin-Film Transistor With a Self-Aligned Etch Stopper Patterned by Back-Side UV Exposure". In: *IEEE Electron Device Letters* 32.6 (2011), pp. 758–760. DOI: 10.1109/LED.2011.2122330.
- [62] Xinwei Geng et al. "How Does Selective Mechanism Improve Self-Attention Networks?" In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 2020, pp. 2986–2995.

- [63] Matteo Ghittorelli et al. “Accurate Analytical Physical Modeling of Amorphous InGaZnO Thin-Film Transistors Accounting for Trapped and Free Charges”. In: *IEEE Transactions on Electron Devices* 61 (2014), pp. 4105–4112. doi: 10.1109/TED.2014.2361062.
- [64] Mostafa Haghi et al. “Wearable Devices in Health Monitoring from the Environmental towards Multiple Domains: A Survey”. In: *Sensors* 21.6 (2021). ISSN: 1424-8220.
- [65] Neska El Haouij et al. “AffectiveROAD system and database to assess driver’s attention”. In: *Proceedings of the 33rd Annual ACM Symposium on Applied Computing*. 2018, pp. 800–803.
- [66] P. Hazucha et al. “Area-efficient linear regulator with ultra-fast load regulation”. In: *IEEE Journal of Solid-State Circuits* 40.4 (2005), pp. 933–940. doi: 10.1109/JSSC.2004.842831.
- [67] Georges Hebrail and Alice Berard. *Individual Household Electric Power Consumption*. UCI Machine Learning Repository. 2006. doi: 10.24432/C58K54.
- [68] Wenzheng Heng, Samuel Solomon, and Wei Gao. “Flexible Electronics and Devices as Human–Machine Interfaces for Medical Robotics”. In: *Advanced Materials* 34.16 (2022), p. 2107902.
- [69] Mei-Chen Hsueh, Timothy K. Tsai, and Ravishankar K. Iyer. “Fault injection techniques and tools”. In: *IEEE Computer* 30.4 (1997), pp. 75–82. doi: 10.1109/2.585157.
- [70] Wei-Hsing Huang et al. “Hardware Acceleration of Kolmogorov-Arnold Network (KAN) for Lightweight Edge Inference”. In: *30th Asia and South Pacific Design Automation Conference (ASPDAC ’25)*. 2025.
- [71] Konstantinos Iordanou et al. “Tiny Classifier Circuits: Evolving Accelerators for Tabular Data”. In: *arXiv:2303.00031* (2023).
- [72] Talha Iqbal et al. “Stress monitoring using wearable sensors: A pilot study and stress-predict dataset”. In: *Sensors* 22.21 (2022), p. 8135.
- [73] A. M. Isaacs et al. “Evaluation of Adult Patients With Hydrocephalus Treated With Endoscopic Third Ventriculostomy With Choroid Plexus Cauterization”. In: *Neurosurgery Focus* 41.3 (2009), E11. doi: 10.3171/2016.8.FOCUS16260.

- [74] Abbas Jamshidi-Roudbari, Po-Chin Kuo, and Miltiadis K. Hatalis. “A flash analog to digital converter on stainless steel foil substrate”. In: *Solid-State Electronics* 54.4 (2010), pp. 410–416. ISSN: 0038-1101. DOI: 10.1016/j.sse.2009.10.020.
- [75] Inhea et al. Jeong. “Machine learning in biosignal analysis from wearable devices”. In: *Materials Horizons* (2025).
- [76] Lakshmikanth Kadge. *Flexible Electronics*. Accessed: 2025-02-28. 2019. URL: <https://edu.ieee.org/in-dscep/2019/12/11/flexible-electronics/>.
- [77] Dervis Karaboga. *An idea based on honey bee swarm for numerical optimization*. Tech. rep. TR06. Erciyes University, Engineering Faculty, Computer Engineering Department, 2005.
- [78] Dervis Karaboga and Bahriye Basturk. “A powerful and efficient algorithm for numerical function optimization: artificial bee colony (ABC) algorithm”. In: *Journal of Global Optimization* 39.3 (2007), pp. 459–471. DOI: 10.1007/s10898-007-9149-x.
- [79] Suhas Krishna Kashyap et al. “Structural Built In Self Test of Analog Circuits using ON/OFF Keying and Delay Monitors”. In: *2024 IEEE 42nd VLSI Test Symposium (VTS)*. IEEE, 2024, pp. 1–7.
- [80] Tejaswini Keragodu et al. “A Voltage Controlled Oscillator Using IGZO Thin-Film Transistors”. In: *2018 IEEE International Symposium on Circuits and Systems (ISCAS)*. 2018, pp. 1–5. DOI: 10.1109/ISCAS.2018.8351175.
- [81] Bongjun Kim et al. “Voltage-Controlled Ring Oscillators Based on Inkjet Printed Carbon Nanotubes and Zinc Tin Oxide”. In: *ACS Applied Materials & Interfaces* 7 (2015). DOI: 10.1021/acsami.5b02093.
- [82] Andreï Kolmogorov. *On the representation of continuous functions of several variables by superpositions of continuous functions of a smaller number of variables*. American Mathematical Society, 1961.
- [83] Andrei N. Kolmogorov. “On the representation of continuous functions of several variables by superpositions of continuous functions of a smaller number of variables”. In: *Doklady Akademii Nauk SSSR* 108 (1957), pp. 179–182.
- [84] Akshatha Kumary and Satheesh Rao. “Design Techniques of Flash ADC: Review”. In: *Advances in Communication, Signal Processing, VLSI, and Embedded Systems*. Springer Singapore, 2020, pp. 89–95.

- [85] Senentxu Lanceros-Méndez and Carlos Miguel Costa. *Printed Batteries: Materials, Technologies and Applications*. Wiley, 2018.
- [86] Jin-Woo Lee et al. “High sensitivity flexible paper temperature sensor and body-attachable patch for thermometers”. In: *Sensors and Actuators A: Physical* 313 (2020), p. 112205. DOI: 10.1016/j.sna.2020.112205.
- [87] Katherine Shu-Min Li et al. “Oscillation ring based interconnect test scheme for SOC”. In: *Proceedings of the Asia and South Pacific Design Automation Conference, ASP-DAC*. Vol. 1. 2005, pp. 184–187.
- [88] Ying-Zu Lin et al. “An Asynchronous Binary-Search ADC Architecture With a Reduced Comparator Count”. In: *IEEE Transactions on Circuits and Systems I: Regular Papers* 57.8 (2010), pp. 1829–1837.
- [89] J. Liu et al. “Recent Progress in Wearable Near-Sensor and In-Sensor Intelligent Perception Systems”. In: *Sensors* 24.7 (2024), p. 2180.
- [90] T. Liu et al. “A built-in self-test technique for load inductance and lossless current sensing of DC-DC converters”. In: *Proceedings of the IEEE 32nd VLSI Test Symposium (VTS-Napa’14)*. CA, 2014, pp. 1–6.
- [91] Yao-Hong Liu et al. “An Ultra-Low Power 1.7–2.7 GHz Fractional-N Sub-Sampling Digital Frequency Synthesizer and Modulator for IoT Applications in 40 nm CMOS”. In: *IEEE Transactions on Circuits and Systems I: Regular Papers* 64.5 (2017), pp. 1094–1105. DOI: 10.1109/TCSI.2016.2625462.
- [92] Ziming Liu et al. “KAN: Kolmogorov-Arnold Networks”. In: *The Thirteenth International Conference on Learning Representations* (2025).
- [93] Stuart P. Lloyd. “Least squares quantization in PCM”. In: *IEEE Transactions on Information Theory* 28.2 (1982), pp. 129–137. DOI: 10.1109/TIT.1982.1056489.
- [94] Paula Carolina Lozano Duarte et al. “Function Approximation Using Analog Building Blocks in Flexible Electronics”. In: *2025 26th International Symposium on Quality Electronic Design (ISQED)*. 2025, pp. 1–7.
- [95] Paula Carolina Lozano Duarte, Sule Ozev, and Mehdi Tahoori. “Distributed Delay-Based BIST for Mixed-Signal Circuits in Flexible Electronics”. In: *31st European Test Symposium (ETS)*. 2026.

- [96] Paula Carolina Lozano Duarte, Sule Ozev, and Mehdi Tahoori. “Fault Tolerant Design of IGZO-based Binary Search ADCs”. In: *2026 27th International Symposium on Quality Electronic Design (ISQED)*. 2026.
- [97] Paula Carolina Lozano Duarte, Sule Ozev, and Mehdi Tahoori. “Spectral Signatures for Parametric Fault Detection in Flexible Electronics”. In: 2026.
- [98] Paula Carolina Lozano Duarte, Georgios Zervakis, and Mehdi Tahoori. “Low-Power PLL-based Clock Stabilization for Flexible IGZO AMS Systems”. In: 2026.
- [99] Paula Carolina Lozano Duarte et al. “Co-Optimization of Analog Kolmogorov-Arnold Networks for Low-Power Function Approximation in Flexible Electronics”. In: 2026.
- [100] Paula Carolina Lozano Duarte et al. “Design and In-training Optimization of Binary Search ADC for Flexible Classifiers”. In: *30th Asia and South Pacific Design Automation Conference (ASPDAC '25)*. 2025. doi: 10.1145/3658617.3697715.
- [101] Paula Carolina Lozano Duarte et al. “Low-Power Binary Search ADC with High-Level Non-Uniform Input Optimization for Flexible Electronics”. In: *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*. 2026.
- [102] Siyuan Lu et al. *Testing and Fault Tolerance Techniques for Carbon Nanotube-Based FPGAs*. 2025. arXiv: 2508.20304 [cs.AR].
- [103] Xiaoyu Luo. “Application of Inkjet Printing Technology in Developing Indicators/Sensors for Intelligent Packaging Systems”. In: *Current Opinion in Food Science* 46 (2022), p. 100868. doi: 10.1016/j.cofs.2022.100868.
- [104] Xiaoyu Luo. “Application of inkjet-printing technology in developing indicators/sensors for intelligent packaging systems”. In: *Current Opinion in Food Science* 46 (2022), p. 100868. doi: 10.1016/j.cofs.2022.100868.
- [105] James MacQueen. “Some methods for classification and analysis of multivariate observations”. In: *Proceedings of the 5th Berkeley Symposium on Mathematical Statistics and Probability*. Vol. 1. University of California Press, 1967, pp. 281–297.

-
- [106] Chris J. Maddison, Andriy Mnih, and Yee Whye Teh. “The Concrete Distribution: A Continuous Relaxation of Discrete Random Variables”. In: *International Conference on Learning Representations (ICLR)*. 2017.
- [107] Mallory Mativenga et al. “Transparent Flexible Circuits Based on Amorphous-Indium–Gallium–Zinc–Oxide Thin-Film Transistors”. In: *IEEE Electron Device Letters* 32.2 (2011), pp. 170–172. DOI: 10.1109/LED.2010.2093504.
- [108] Joel Max. “Quantizing for minimum distortion”. In: *IRE Transactions on Information Theory* 6.1 (1960), pp. 7–12. DOI: 10.1109/TIT.1960.1057548.
- [109] Carver Mead. *Analog VLSI and neural systems*. USA: Addison-Wesley Longman Publishing Co., Inc., 1989.
- [110] Zikang Mei et al. “Advances in Organic Thin-Film Transistor Technology for Flexible Analog Front Ends in Wearable Electronics”. In: *IEEE Journal on Flexible Electronics* 3.11 (2024), pp. 508–521. DOI: 10.1109/JFLEX.2025.3528904.
- [111] Fanzhao Meng et al. “A Performance Optimized Operational Amplifier Using Transconductance Enhancement Topology Based on a-IGZO TFTs”. In: *IEEE Journal of the Electron Devices Society* 12 (2024), pp. 159–164. DOI: 10.1109/JEDS.2024.3366554.
- [112] Pooria Mostafalu et al. “Wireless Flexible Smart Bandage for Continuous Monitoring of Wound Oxygenation”. In: *IEEE Transactions on Biomedical Circuits and Systems* 9.5 (2015), pp. 670–677. DOI: 10.1109/TBCAS.2015.2488582.
- [113] Vojtech Mrazek et al. “Arbitrary Precision Printed Ternary Neural Networks With Holistic Evolutionary Approximation”. In: *IEEE Transactions on Circuits and Systems for Artificial Intelligence* 2.4 (2025), pp. 351–363.
- [114] Vojtech Mrazek et al. “Design of power-efficient approximate multipliers for approximate artificial neural networks”. In: *Int. Conf. on Computer-Aided Design (ICCAD)*. 2016.
- [115] Muhammad Husnain Mubarak et al. “Printed Machine Learning Classifiers”. In: *Annu. Int. Symp. Microarchitecture (MICRO)*. 2020, pp. 73–87. DOI: 10.1109/MICRO50266.2020.00019.

- [116] Amlan Nag and Roy P Pailly. “Low power squaring and square root circuits using subthreshold MOS transistors”. In: *2009 International Conference on Emerging Trends in Electronic and Photonic Devices & Systems*. IEEE. 2009, pp. 96–99.
- [117] Arokia Nathan et al. “Flexible Electronics: The Next Ubiquitous Platform”. In: *Proceedings of the IEEE 100*. Special Centennial Issue (2012), pp. 1486–1517.
- [118] T. Oshita et al. “A compact first-order $\Sigma\Delta$ modulator for analog high-volume testing of complex system-on-chips in a 14nm tri-gate digital CMOS”. In: *IEEE Journal of Solid-State Circuits* 51.2 (2016).
- [119] Emre Ozer et al. “Bespoke Machine Learning Processor Development Framework on Flexible Substrates”. In: *Int. Conf. Flexible and Printable Sensors and Systems (FLEPS)*. 2019, pp. 1–3.
- [120] P. Pal et al. “Special Session: Reliable Emerging Electronics in Wearable and Implantable Healthcare Applications”. In: *2026 IEEE VLSI Test Symposium (VTS)*. 2026.
- [121] Priyanjana Pal et al. “Analog Printed Spiking Neuromorphic Circuit”. In: *IEEE DATE*. 27th Design, Automation and Test in Europe Conference (Valencia, Spain, Mar. 25–27, 2024). 2024.
- [122] Zhong Pan et al. “Approaches to Improve Mobility and Stability of IGZO TFTs: A Brief Review”. In: *Transactions on Electrical and Electronic Materials* 25 (2024). DOI: 10.1007/s42341-024-00536-1.
- [123] Konstantin L Panchuk, Tatyana Mikhailovna Myasoedova, and Mariya N Odinets. “Construction of a discrete planar contour by fractional rational Bezier curves of second order”. In: *Journal of Physics: Conference Series*. Vol. 1441. 1. IOP Publishing. 2020, p. 012072.
- [124] Nikolas P. Papadopoulos et al. “Toward Temperature Tracking With Unipolar Metal-Oxide Thin-Film SAR C-2C ADC on Plastic”. In: *IEEE Journal of Solid-State Circuits* 53.8 (2018), pp. 2263–2272.
- [125] PragmatIC Semiconductor. *FlexICs - Flexible Integrated Circuits*. Accessed: 15 March 2025. 2025. URL: <https://www.pragmaticsemi.com/flexics>.
- [126] Alireza Raisiardi et al. “Flexing RISC-V Instruction Subset Processors to the Extreme Edge”. In: *Proceedings of the 58th IEEE/ACM International Symposium on Microarchitecture (MICRO)*. 2025.

- [127] Daniele Raiteri et al. “An organic VCO-based ADC for quasi-static signals achieving 1LSB INL at 6b resolution”. In: *2013 IEEE International Solid-State Circuits Conference Digest of Technical Papers*. 2013, pp. 108–109.
- [128] Farhan Rasheed et al. “Variability Modeling for Printed Inorganic Electrolyte-Gated Transistors and Circuits”. In: *IEEE Transactions on Electron Devices* 66.1 (2019), pp. 146–152.
- [129] Behzad Razavi. *Design of CMOS Phase-Locked Loops: From Circuit Level to Architecture Level*. Cambridge University Press, 2020. doi: 10.1017/9781108626200.
- [130] Attila Reiss, Ina Indlekofer, and Philip Schmidt. *PPG-DaLiA*. UCI Machine Learning Repository. 2019. doi: 10.24432/C53890.
- [131] Harold L. Rekate. “A contemporary definition and classification of hydrocephalus”. In: *Seminars in Pediatric Neurology* 16.1 (2011), pp. 9–15. doi: 10.1016/j.spen.2009.01.002.
- [132] Gabriel A. Rincon-Mora and Phillip E. Allen. “A low-voltage, low quiescent current, low drop-out regulator”. In: *IEEE Journal of Solid-State Circuits* 33.1 (1998), pp. 36–44. doi: 10.1109/4.654935.
- [133] *RMS-to-DC Conversion Techniques*. Application Note AN-691. Analog Devices. 2017.
- [134] G. W. Roberts, F. Taenzler, and M. Burns. *An Introduction to Mixed-Signal IC Test and Measurement*. Oxford University Press, 2011.
- [135] Rahul Sarpeshkar. “Analog Versus Digital: Extrapolating from Electronics to Neurobiology”. In: *Neural Computation* 10.7 (1998), pp. 1601–1638.
- [136] Philip Schmidt et al. “Introducing WESAD, a Multimodal Dataset for Wearable Stress and Affect Detection”. In: *Proceedings of the 20th ACM International Conference on Multimodal Interaction*. Association for Computing Machinery, 2018, pp. 400–408.
- [137] Adel S. Sedra and Kenneth C. Smith. *Microelectronic Circuits*. 7th ed. Oxford University Press, 2015.
- [138] Lukas Sekanina, Zdenek Vasicek, and Vojtech Mrazek. “Automated Search-Based Functional Approximation for Digital Circuits”. In: *Approximate Circuits: Methodologies and CAD*. Ed. by Sherief Reda and Muhammad Shafique. Springer International Publishing, 2019, pp. 175–203. doi: 10.1007/978-3-319-99322-5_9.

- [139] Maha Shatta et al. “Feature-to-Classifer Co-Design for Mixed-Signal Smart Flexible Wearables for Healthcare at the Extreme Edge”. In: *Proceedings of the IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*. 2025.
- [140] SS Sidharth et al. *Chebyshev Polynomial-Based Kolmogorov-Arnold Networks: An Efficient Architecture for Nonlinear Function Approximation*. 2024. arXiv: 2405.07200 [cs.LG].
- [141] H. Spence. “Automatic analog fault simulation”. In: *Conference Record. AUTOTESTCON '96*. 1996, pp. 17–22. DOI: 10.1109/AUTEST.1996.547671.
- [142] T. Stuart, J. Hanna, and P. Gutruf. “Wearable devices for continuous monitoring of biosignals: Challenges and opportunities”. In: *APL Bioengineering* 6.2 (2022), p. 021502. DOI: 10.1063/5.0086935.
- [143] Arun Suresh et al. “Fast All-Transparent Integrated Circuits Based on Indium Gallium Zinc Oxide Thin-Film Transistors”. In: *IEEE Electron Device Letters* 31.4 (2010), pp. 317–319. DOI: 10.1109/LED.2010.2041525.
- [144] Sun T. et al. “Artificial Intelligence Meets Flexible Sensors: Emerging Smart Flexible Sensing Systems Driven by Machine Learning and Artificial Synapses”. In: *Nano-Micro Letters* (2024). DOI: 10.1007/s40820-023-01235-x.
- [145] Mehdi B Tahoori et al. “Computing with printed and flexible electronics”. In: *2025 IEEE European Test Symposium (ETS)*. IEEE. 2025, pp. 1–9.
- [146] Xiyuan Tang et al. “Low-Power SAR ADC Design: Overview and Survey of State-of-the-Art Techniques”. In: *IEEE Transactions on Circuits and Systems I: Regular Papers* 69.6 (2022), pp. 2249–2262. DOI: 10.1109/TCSI.2022.3166792.
- [147] Koichi Tanno, Okihiko Ishizuka, and Zheng Tang. “Four-quadrant CMOS current-mode multiplier independent of device parameters”. In: *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing* 47.5 (2000), pp. 473–477.

- [148] Javad Tavakoli, Hossein Miri Lavasani, and Samad Sheikhaei. “An Ultra Low Power Integer-N PLL with a High-Gain Sampling Phase Detector for IoT Applications in 65 nm CMOS”. In: *Journal of Low Power Electronics and Applications* 13.4 (2023), p. 65. DOI: 10.3390/jlpea13040065.
- [149] M. F. Toner and G. W. Roberts. “A BIST scheme for an SNR test of a sigma-delta ADC”. In: *Proceedings of the IEEE International Test Conference*. Washington, DC, 1993, pp. 805–814.
- [150] Duy Tran Van et al. “Exploring the Limitations of Kolmogorov-Arnold Networks in Classification: Insights to Software Training and Hardware Implementation”. In: 2024. arXiv: 2407.17790.
- [151] Zdenek Vasicek and Lukas Sekanina. “Evolutionary Approach to Approximate Digital Circuits Design”. In: *IEEE Transactions on Evolutionary Computation* 19.3 (2015), pp. 432–444. DOI: 10.1109/TEVC.2014.2336175.
- [152] Theofanis Vergos et al. “Bespoke Co-Processor for Energy-Efficient Health Monitoring on RISC-V-Based Flexible Wearables”. In: *Proceedings of the Design, Automation and Test in Europe Conference (DATE)*. 2026.
- [153] Nishtha Wadhwa et al. “Mostly Passive Δ - Σ ADC with a-IGZO TFTs for Flexible Electronics”. In: *2021 IEEE International Symposium on Circuits and Systems (ISCAS)*. 2021, pp. 1–6.
- [154] Dennis D Weller et al. “Realization and Training of an Inverter-based Printed Neuromorphic Computing System”. In: *Scientific Reports* 11.1 (2021), pp. 1–13.
- [155] Dennis D. Weller et al. “Printed Stochastic Computing Neural Networks”. In: *Design Automation and Test in Europe (DATE)*. 2021, pp. 914–919.
- [156] Zhu Ying et al. “Indium–gallium–zinc–oxide thin-film transistors: Materials, devices, and applications”. In: *Journal of Semiconductors* 42.3 (2021), p. 031101.
- [157] Shengqi Yu et al. “Approximate digital-in analog-out multiplier with asymmetric nonvolatility and low energy consumption”. In: *Microelectronics Journal* (2023). DOI: 10.1016/j.vlsi.2023.05.009.

- [158] Georgios Zervakis et al. “VADER: Voltage-Driven Netlist Pruning for Cross-Layer Approximate Arithmetic Circuits”. In: *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.* 27.6 (2019), pp. 1460–1464. DOI: 10 . 1109/TVLSI.2019.2900160.
- [159] Georgios Zervakis, Hussam Amrouch, and Jörg Henkel. “Design Automation of Approximate Circuits With Runtime Reconfigurable Accuracy”. In: *IEEE Access* 8 (2020), pp. 53522–53538. DOI: 10 . 1109 / ACCESS.2020.2981395.
- [160] An Qi Zhang et al. “Stuck-at Faults Tolerance and Recovery in MLP Neural Networks Using Imperfect Emerging CNFET Technology”. In: *IEEE Journal on Exploratory Solid-State Computational Devices and Circuits* 9.2 (2023), pp. 168–175. DOI: 10.1109/JXCDC.2023.3313127.
- [161] Haibin Zhao et al. “Aging-Aware Training for Printed Neuromorphic Circuits”. In: *IEEE/ACM International Conference On Computer Aided Design (ICCAD)*. 2022, pp. 1–9.
- [162] Haibin Zhao et al. “Highly-Bespoke Robust Printed Neuromorphic Circuits”. In: *Design, Automation & Test in Europe Conference & Exhibition (DATE)*. 2023. DOI: 10.23919/DATE56975.2023.10137298.
- [163] Haibin Zhao et al. “Power-Aware Training for Energy-Efficient Printed Neuromorphic Circuits”. In: *International Conference on Computer Aided Design (ICCAD)*. 2023, pp. 1–9. DOI: 10 . 1109 / ICCAD57390.2023.10323917.
- [164] Haibin Zhao et al. “Printed Electrodermal Activity Sensor with Optimized Filter for Stress Detection”. In: *ISWC’22, Atlanta, GA and Cambridge, UK, September 11–15, 2022*.
- [165] Haibin Zhao et al. “Towards Temporal Information Processing – Printed Neuromorphic Circuits with Learnable Filters”. In: *Proceedings of the 18th ACM International Symposium on Nanoscale Architectures*. Association for Computing Machinery, 2024.
- [166] Kemeng Zhou et al. “Printable and flexible integrated sensing systems for wireless healthcare”. In: *Nanoscale* 16.15 (2024), pp. 7264–7286. DOI: 10.1039/D3NR06099C.
- [167] Johanna Zikulnig et al. “Printed Electronics Technologies for Additive Manufacturing of Hybrid Electronic Sensor Systems”. In: *Advanced Sensor Research* 2.7 (2023), p. 2200073.

- [168] Barret Zoph and Quoc V. Le. “Neural Architecture Search with Reinforcement Learning”. In: *International Conference on Learning Representations (ICLR)*. 2017.
- [169] Mohammad Zulqarnain, Stefano Stanzione, and Eugenio Cantatore. “A Digitally Assisted Tunable High Pass Filter Based on Flexible a-IGZO TFTs for Biomedical Applications”. In: *2020 IEEE International Symposium on Circuits and Systems (ISCAS)*. 2020, pp. 1–5. DOI: 10.1109/ISCAS45731.2020.9180656.
- [170] Christoph Zysset et al. “IGZO TFT-Based All-Enhancement Operational Amplifier Bent to a Radius of 5 mm”. In: *IEEE Electron Device Letters* 34.11 (2013), pp. 1394–1396. DOI: 10.1109/LED.2013.2280024.

A. Acronyms and Symbols

ABB	Analog Building Block
ABC	Artificial Bee Colony
ADC	Analog-to-Digital Converter
AKAN	Analog Kolmogorov-Arnold Network
AMS	Analog and Mixed-Signal
ATE	Automatic Test Equipment
AxC	Approximate Computing
BIST	Built-In Self-Test
BVP	Blood Volume Pulse
CGP	Cartesian Genetic Programming
CMOS	Complementary Metal-Oxide-Semiconductor
CP	Charge Pump
CSF	Cerebrospinal Fluid
CUT	Circuit Under Test
DAC	Digital-to-Analog Converter
DISR	Double Input Symmetrical Relevance
DNAS	Differentiable Neural Architecture Search
DNL	Differential Nonlinearity
DSE	Design Space Exploration
DT	Decision Tree
EDA	Electrodermal Activity
EGFET	Electrolyte-Gated Field-Effect Transistor
ESR	Equivalent Series Resistance
FE	Flexible Electronics
FlexIC	Flexible Integrated Circuit
GSR	Galvanic Skin Response
HPF	High-Pass Filter
IGZO	Indium Gallium Zinc Oxide
IoT	Internet of Things
ISM	Industrial, Scientific and Medical
JMI	Joint Mutual Information
KAN	Kolmogorov-Arnold Network
LDO	Low-Dropout Regulator
LPF	Low-Pass Filter
LTG	Linear Threshold Gate

LTP	Lottery Ticket Pruning
LTPS	Low-Temperature Polycrystalline Silicon
MLP	Multi-Layer Perceptron
MUX	Multiplexer
NAS	Neural Architecture Search
NMSE	Normalized Mean Square Error
NSGA	Non-dominated Sorting Genetic Algorithm
PDK	Process Design Kit
PFD	Phase-Frequency Detector
PLL	Phase-Locked Loop
PSS	Periodic Steady State
PSRR	Power Supply Rejection Ratio
PVT	Process-Voltage-Temperature
RF	Radio Frequency
RMS	Root Mean Square
RO	Ring Oscillator
RTL	Resistor-Transistor Logic
SAR	Successive Approximation Register
SVM	Support Vector Machine
TEMP	Skin Temperature
TFT	Thin-Film Transistor
TMR	Triple Modular Redundancy
TNN	Ternary Neural Network
VCO	Voltage-Controlled Oscillator
VLSI	Very Large Scale Integration
WESAD	Wearable Stress and Affect Detection

B. List of Publications

First-Author Publications

1. **P. C. Lozano Duarte**, A. Dube, G. Zervakis, M. Tahoori, S. Nassif. "Function Approximation Using Analog Building Blocks in Flexible Electronics." *26th International Symposium on Quality Electronic Design (ISQED)*, 2025.
2. **P. C. Lozano Duarte**, G. Zervakis, M. Tahoori, S. Nassif. "Co-Optimization of Analog Kolmogorov-Arnold Networks for Low-Power Function Approximation in Flexible Electronics." *IEEE Journal on Emerging and Selected Topics in Circuits and Systems (IEEE JETCAS)*, 2026.
3. **P. C. Lozano Duarte**, F. Afentaki, G. Zervakis, M. Tahoori. "Design and In-Training Optimization of Binary Search ADC for Flexible Classifiers." *30th Asia and South Pacific Design Automation Conference (ASP-DAC)*, 2025.
4. **P. C. Lozano Duarte**, F. Afentaki, G. Zervakis, M. Tahoori. "Low-Power Binary Search ADC with High-Level Non-Uniform Input Optimization for Flexible Electronics." *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (IEEE TCAD)*.
5. **P. C. Lozano Duarte**, S. Ozev, M. Tahoori. "Fault Tolerant Design of IGZO-based Binary Search ADCs." *27th International Symposium on Quality Electronic Design (ISQED)*, 2026.
6. **P. C. Lozano Duarte**, S. Ozev, M. Tahoori. "Distributed Delay-Based BIST for Mixed-Signal Circuits in Flexible Electronics." *European Test Symposium (ETS)*, 2026.
7. **P. C. Lozano Duarte**, S. Ozev, M. Tahoori. "Spectral Signatures for Parametric Fault Detection in Flexible Electronics."
8. **P. C. Lozano Duarte**, G. Zervakis, M. Tahoori. "Low-Power PLL-Based Clock Stabilization for Flexible IGZO AMS Systems."

Co-Authored Publications

1. V. Mrazek, K. Balaskas, **P. C. Lozano Duarte**, Z. Vasicek, M. Tahoori, G. Zervakis. “Arbitrary Precision Printed Ternary Neural Networks with Holistic Evolutionary Approximation.” *IEEE Transactions on Circuits and Systems for Artificial Intelligence (TCAS-AI)*, 2025.
2. G. Armeniakos, **P. C. Lozano Duarte**, P. Pal, G. Zervakis, M. Tahoori, D. Soudris. “On-Sensor Printed ML Classification via Bespoke ADC and Decision Tree Co-Design.” *Design, Automation & Test in Europe (DATE)*, 2024.
3. F. Afentaki, **P. C. Lozano Duarte**, G. Zervakis, M. Tahoori. “Reducing ADC Front-End Costs during Training of On-Sensor Printed Multilayer Perceptrons.” *IEEE Embedded Systems Letters*, vol. 16, no. 4, 2024.
4. P. Pal, **P. C. Lozano Duarte**, S. K. Kashyap et al. “Special Session: Reliable Emerging Electronics in Wearable and Implantable Healthcare Applications.” *IEEE VLSI Test Symposium (VTS) Special Session*, 2026.
5. F. Afentaki, **P. C. Lozano Duarte**, G. Zervakis, M. Tahoori. “Lightweight Fault Resilient Flexible Flash-ADCs.” *International Symposium on On-Line Testing and Robust System Design (IOLTS)*, 2026.
6. F. Afentaki, S.S.R. Nakkilla, K. Balaskas, **P. C. Lozano Duarte** et al. “Exploration of Low-Power Flexible Stress Monitoring Classifiers for Conformal Wearables.” *International Symposium on Low Power Electronics and Design (ISLPED)*, 2025.
7. F. Afentaki, K. Balaskas, A. Datta, **P. C. Lozano Duarte** et al. “Real-Time Stress Monitoring with Flexible Electronics: A Co-Design Approach from ADCs to Feature Extractors and Classifiers.” *IEEE Internet of Things Journal (IEEE IoT)*, 2026.
8. M. Shatta, K. Balaskas, **P. C. Lozano Duarte** et al. “Feature-to-Classifier Co-Design for Mixed-Signal Smart Flexible Wearables for Healthcare at the Extreme Edge.” *IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, 2025.