



Delft University of Technology

Document Version

Final published version

Licence

CC BY

Citation (APA)

Strydis, C., Veselka, D., van der Kleij, J., & Siddiqi, M. A. (2026). FPGA-Enabled, Smart Implants: A Path to Personalized Healthcare. *IEEE Open Journal of Engineering in Medicine and Biology*, 7, 291-300.
<https://doi.org/10.1109/OJEMB.2026.3720369>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.

FPGA-Enabled, Smart Implants: A Path to Personalized Healthcare

Christos Strydis , Senior Member, IEEE, David Veselka, Job van der Kleij, and Muhammad Ali Siddiqi 

Abstract—Goal: Contrary to the conventional view that Field-Programmable Gate Arrays (FPGAs) are poorly suited to the stringent energy and resource constraints of Implantable Medical Devices (IMDs), this work investigates their potential to improve computational performance and energy efficiency in this domain. We propose heterogeneous CPU–FPGA systems as a means to support the growing demand for embedded AI in next-generation IMDs for intelligent and personalized healthcare. **Methods:** We evaluated heterogeneous CPU–FPGA architectures against CPU-only configurations across representative AI-driven neural and security IMD workloads. **Results:** FPGAs achieved lower execution latency and improved energy efficiency than CPU-based implementations in computationally intensive neural workloads, despite higher instantaneous power consumption in some cases due to parallel execution. In low-duty-cycle security workloads where aggressive power gating is feasible, FPGA- and CPU-based solutions exhibited comparable overall performance. **Conclusions:** These findings demonstrate that heterogeneous CPU–FPGA systems can integrate advanced AI and security algorithms into IMDs while satisfying battery-life, power-density, real-time processing, and practical implantability constraints relevant to long-term medical-device deployment, enabling intelligent, personalized, and secure implantable healthcare.

Index Terms—Energy efficiency, embedded AI, field-programmable gate array, heterogeneous system, implantable medical device.

Received 4 November 2025; revised 4 June 2026; accepted 26 July 2026. Date of publication 4 August 2026; date of current version 21 August 2026. This work was supported in part by the European Union's Horizon Europe R&I programme through projects SEPTON and SECURED under Grant 101094901 and Grant 101095717, in part by the Dutch Research Council (NWO) through the Gravitation Programme DBI under Grant 024.005.022, and in part by the Lahore University of Management Sciences through the Faculty Initiative Fund under Grant FIF-0876, Grant FIF-0961, and Grant FIF-1047. The review of this article was arranged by Editor Thomas Stieglitz. (Corresponding author: Christos Strydis.)

Christos Strydis is with Neuroscience Department, Erasmus Medical Center, 3015 GD Rotterdam, The Netherlands, and also with the Quantum & Computer Engineering Department, Delft University of Technology, 2628 CD Delft, The Netherlands (e-mail: c.strydis@tudelft.nl).

David Veselka and Job van der Kleij are with Quantum and Computer Engineering Department, Delft University of Technology, 2628 CD Delft, The Netherlands.

Muhammad Ali Siddiqi is with the Electrical Engineering Department, Lahore University of Management Sciences, Lahore 54792, Pakistan, also with the Quantum and Computer Engineering Department, Delft University of Technology, 2628 CD, Delft, The Netherlands, and also with the Neuroscience Department, Erasmus Medical Center, 3015 GD Rotterdam, The Netherlands (e-mail: m.siddiqi@lums.edu.pk).

Digital Object Identifier 10.1109/OJEMB.2026.3720369

Impact Statement—FPGAs can substantially enhance computational performance in future implantable medical devices delivering embedded AI and security while preserving energy efficiency across diverse operational regimes.

I. INTRODUCTION

IMPLANTABLE Medical Devices (IMDs) are rapidly evolving from simple, low-power systems into intelligent platforms capable of executing emerging artificial intelligence (AI)/machine learning (ML)-oriented processing workloads for advanced neural interventions [1], [2], [3]. Recent advances in AI have enabled real-time interpretation of complex neural signals [4], [5], promising patient-specific improvements in neurotherapeutic outcomes. Yet, as implants increasingly communicate via mainstream wireless protocols such as Bluetooth LE, traditional IMD designs, optimized for static, low-power operation, are facing unprecedented demands in both computing and security [3], [6].

This paradigm shift introduces two key challenges. First, emerging AI/ML-oriented functionalities raise computational demands inside implants, requiring architectures that execute complex algorithms with low latency and support in vivo updates, from security patches to adaptive control. Second, wireless connectivity exposes implants to privacy and cybersecurity risks that conventional designs cannot effectively mitigate [6], [7]. These constraints also introduce broader safety, reliability, and long-term deployment considerations relevant to practical medical-device regulation and implantability.

We propose addressing these challenges through the integration of reconfigurable-fabric logic, namely tiny Field-Programmable Gate Arrays (FPGAs), within future IMDs (Fig. 1). Despite skepticism regarding their suitability for ultra-low-power environments, FPGAs offer real-time, energy-efficient computation and post-deployment reconfigurability, transforming implants from static medical tools into adaptive, secure computing platforms. However, questions remain regarding their energy efficiency, latency impact, and compliance with stringent medical regulations.

Early explorations of reconfigurable logic in implants [8], [9] demonstrated proof-of-concept designs but did not reach clinical translation. Instead, modern implants predominantly employ Microcontroller Units (MCUs) or CPUs, sometimes augmented with dedicated ASIC blocks for cryptographic tasks (e.g., AES engines) [10]. While ASICs offer excellent power efficiency, their rigidity limits adaptability to evolving medical




Fig. 1. Top: An example IMD (deep-brain neurostimulator). Bottom: Major energy consumers in a typical IMD. The dotted boxes indicate the blocks added to study FPGA feasibility.

and security requirements. Besides, GPUs, though highly parallel, are often too power-hungry for IMDs. In contrast, FPGAs, though typically more power-hungry, can be reprogrammed post-implantation, support algorithmic evolution, and remain cost-effective for low-to-medium device volumes.

Recent research has begun to re-evaluate FPGAs' role in energy-constrained biomedical systems. Khan and da Silva [11] review FPGA applications in wearable devices, identifying low-power families and techniques such as sleep modes, voltage scaling, partial reconfiguration, and hybrid flash architectures: approaches directly relevant to IMDs. Similarly, Vaithianathan et al. [12] outline strategies for optimizing FPGAs for implantable and wearable devices, discussing components and design principles to enhance efficiency, though without empirical validation. Altman et al. [13] provide a broader review of FPGA applications in medical machine learning, highlighting the growing role of reconfigurable hardware in biomedical computation.

Despite these developments, FPGA-enabled implants remain underexplored experimentally. In particular, a systematic assessment jointly considering AI-driven neural workloads, security processing, physiological duty-cycle constraints, FPGA reconfiguration overheads, battery-life implications, and regulatory considerations in the context of next-generation secure and intelligent implants is still missing. In this work, we examine whether FPGAs can realistically serve as viable building blocks for such systems by evaluating their performance in energy-constrained conditions and analyzing their implications for future regulatory and design frameworks. Our findings aim to chart a path toward

adaptive, updatable, and secure medical implants that bridge the gap between computational neuroscience, embedded systems, and cybersecurity. Concisely, the contributions of this work from a system-level feasibility and deployment perspective are as follows:

- 1) A systematic comparison of MCU-only and heterogeneous MCU-FPGA IMD architectures across representative neural, cardiac, and security workloads under realistic implant constraints.
- 2) An analysis of the interaction between physiological duty cycles, FPGA power-gating opportunities, and SRAM-based FPGA reconfiguration overheads in implantable systems.
- 3) A realistic assessment of battery-life, power-density, and implantability implications across representative IMD operating scenarios.
- 4) A discussion of deployment and regulatory implications of reconfigurable logic in IMDs under existing medical-device frameworks.

II. MATERIALS AND METHODS

The feasibility of employing FPGAs as accelerators in future implants depends critically on the design assumptions and the representativeness of the evaluated workloads. This section outlines the adopted methodology, study parameters, and boundary conditions, together with the selection criteria for processing devices, benchmarks, and workloads. It also details the measurement setup and performance metrics used, establishing a consistent and reproducible framework for evaluation.

A. Boundary Conditions

To ensure a comprehensive evaluation and maximize the experimental realism, we selected two contrasting use cases: a minimal-workload cardiac IMD (e.g., a pacemaker) and a computationally intensive neural IMD (neurostimulator) employing a neural-network-based algorithm. Furthermore, neurostimulators face inherently greater demands due to the stringent real-time constraints imposed by rapidly changing neural signals, in contrast to the comparatively slower cardiac signals processed by pacemakers. Fig. 2 demonstrates that the average inter-spike interval of a Purkinje-cell signal in the cerebellum is orders of magnitude shorter than that of a cardiac signal.

B. Device Choice

The choice of processing devices is critical for evaluating next-generation IMD workloads. In this study, we focus on heterogeneous MCU-FPGA architectures, comparing conventional MCU-based implementations against FPGA-accelerated solutions for representative medical and security workloads.

1) CPU: MCUs remain the baseline processing units for IMDs due to their integration of CPUs, memory, and peripherals. For this study, the Tiny Gecko 11 MCU (EFM32TG11) from Silicon Labs was selected that has the Cortex-M0+ processor from ARM [10]. This device is optimized for low-power applications and includes dedicated hardware for AES encryption, providing a useful comparison for FPGA implementations.




Fig. 2. Physiological signals and the associated time scales. *Top:* A typical cardiac signal. *Bottom:* A (single-unit) neural signal recorded from the cerebellum.

TABLE I
DEVICE CHARACTERISTICS

Device Name	Freq. range used in study (MHz)	Voltage (V)	Current-Measurement Method
Tiny Gecko 11	0.001–48	1.8–3.3	Advance Energy Monitoring via IDE
iCE40UP5K	0.01–48	1.2	Manually via test points 11 & 12 (VCC)

2) FPGA: The Lattice iCE40UP5K was selected as the primary FPGA [14] due to its low-power characteristics and reconfigurability, critical for IMD use cases (see Table I). The reported frequency range corresponds to the operating conditions employed in the experiments and should not be interpreted as the architectural maximum frequency capability of the FPGA fabric.

C. Assumptions and Constraints

To ensure consistency, several realistic assumptions were made:

- **Data-Transfer Rate and Volume:** A transfer rate of 265 kbps was assumed, reflecting the maximum rate employed by a commercial implantable-grade transceiver [15]. Best-case and worst-case transfer durations were calculated for data volumes ranging from 1 MB *monthly* transmissions to 10 MB *daily* transmissions. This range encompasses typical patient accesses using their smartphones and more detailed data transfers associated with readers/programmers of specialists.
- **Duty Cycling:** IMDs are assumed to operate in cycles of active and sleep modes to conserve energy. Each workload is modeled as a periodic task consisting of an active execution interval followed by an idle interval. Duty cycle is defined as the ratio between active execution time and total workload period. The duty cycling percentages are dependent on the workload type, implant type (cardiac or neural), and the data-transfer volume (discussed above). This analysis considers both clock gating (CG) and power gating (PG) as FPGA sleep mode strategies. For the EFM32TG11, ‘power gating’ refers to the deepest low-power sleep mode employed in this study. For the

TABLE II
PACEMAKER AND NEUROSTIMULATOR SPECIFICATIONS

Parameter	Value
Supply voltage	1.8 V (MCU), 1.2 V (FPGA)
Clock frequency	12 MHz
RF-transceiver effective data rate	265 kbps (maximum value) [15]
Battery capacity	6.3 Ah [20]
Pacemaker:	
Stimulation energy	20 μ J per pacing pulse [21]
Average no. of pacing events / day	10000 [22]
Sampling rate	1000 Hz [23]
Neurostimulator:	
Stimulation current	12 mA [24]
Pulse width	1 ms [24]
Pulse freq.	333 Hz [24]
Burst duration	10 seconds [24]
Average seizures / day	4.3 [25]
Sampling rate	24414 Hz [17]

iCE40UP, wake-up from PG requires full SRAM configuration reload from external SPI flash memory, introducing millisecond-scale reconfiguration latency.

- **Reconfiguration Costs:** The analysis accounted for the energy and time overhead of FPGA reconfiguration, particularly during wake-up from power-gated sleep states. While the iCE40UP5K includes both SRAM and Non-volatile Configuration Memory (NVCM) for the active configuration (bitstream), the *one-time* programmability of NVCM necessitates the use of an external SPI Flash (MX25V1006F [16]) for non-volatile bitstream storage. Consequently, FPGA wake-up from PG requires full bitstream reload from external flash memory.
- **Environmental Context:** Section II-A highlights the disparity in computational complexity between cardiac and neural environments, with the latter demanding significantly greater processing power. This difference extends to stimulation strategies, where the unique patterns required for neural modulation result in differing stimulation current consumption compared to cardiac pacing (see Table II). The parameters summarized in the table were used as system-level assumptions for workload scheduling, duty-cycle estimation, stimulation activity, communication activity, and battery-life analysis throughout the evaluated scenarios. These operating assumptions were derived from published literature, commercial-device specifications, and prior implant-research experience [7].

D. Workloads

To evaluate the devices, a series of workloads and scenarios were defined. These workloads encompass both medical and security algorithms, designed to represent real-world use cases for IMDs. An *iteration* refers to one complete execution of the target workload for a single processing event or processing window, depending on the workload type. Two categories of workloads were employed:

1) Medical Algorithms: Next-generation neural workloads, exemplified by spike sorting [17] and Current-Source Density (CSD) [18] algorithms, present significant computational challenges. Spike sorting, encompassing spike *detection* and *classification*, was chosen for its significant computational demands and direct applicability to neural IMDs. We selected spike sorting of Purkinje-cell neural signals (single-unit recordings) from the cerebellum, a process particularly relevant for next-generation implants designed to address brain injuries and motor function deficits. Furthermore, spike classification, performed using a multilayer perceptron (MLP), serves as a representative, embedded AI/ML workload with nontrivial computational demands and limited gating opportunities.

CSD [18], a less computationally intensive algorithm, was included for comparative analysis due to its utility in detecting events within Local-Field Potential (LFP) signals, rather than the single-unit signals previously discussed. Similar to single-unit signals, LFP signals are a critical data type recorded by neural IMDs.

Additionally, spike detection [19] was employed to represent simpler cardiac workloads. In this context, it is used to detect heartbeats, allowing for heart rate computation (using a basic counter) for pacing management.

2) Security Algorithms: Given the critical function of IMDs and the sensitive patient data they handle, robust security primitives are essential. To benchmark security algorithms, we implemented AES-128 (block cipher) to serve as the mainstream primitive, alongside lightweight alternatives such as SIMON/SPECK (block cipher) and PHOTON (hash function), to evaluate energy-efficient cryptographic workloads. Given that numerous modern MCUs, such as the EFM32TG11, incorporate cryptographic accelerators for AES-128, we also included hardware AES (HW AES) measurements from the EFM32TG11 for comparative analysis.

The disparate duty-cycling demands of medical and security workloads, coupled with their independent nature and the iCE40UP's inherent limitations regarding partial reconfiguration, necessitate the use of two FPGAs, each dedicated to its respective workload, as illustrated in Fig. 1.

E. Scenarios and Processing-Device Configurations

Four scenarios were developed to evaluate the devices under varying conditions:

- 1) *Cardiac* IMD with a data-transfer volume of 1-10 MB per day.
- 2) *Cardiac* IMD with a data-transfer volume of 1-10 MB per month.
- 3) *Neural* IMD with a data-transfer volume of 1-10 MB per day.
- 4) *Neural* IMD with a data-transfer volume of 1-10 MB per month.

The following processing-device configurations were evaluated across the defined scenarios:

- 1) *MCU only:* Both the medical and security workloads are implemented as software on EFM32TG11, which enters




Fig. 3. Scenarios assumed in this study based on IMD/workload type and the volume of wireless data transmissions.

power gating during idle periods. Due to the stringent real-time constraints of neural signals, power gating is limited to cardiac spike detection and security workloads, and cannot be applied to neural spike sorting.

- 2) *MCU only (HW AES):* Same as above except that it utilizes EFM32TG11's hardware-accelerated AES-128, rather than a software implementation.
- 3) *MCU (Med) + FPGA (AES):* This is a hybrid solution in which medical workloads are executed on the MCU and security workloads on the FPGA. Power gating is applied to both the devices (but excluding neural spike sorting, as discussed above).
- 4) *FPGA only:* Both the medical and security workloads are executed on the respective FPGAs (see Fig. 1).
- 5) *FPGA only (CG):* Same as *FPGA only* but with clock gating for power savings during inactivity. Because of the strict real-time requirements of neural signals, clock gating is only feasible for cardiac spike detection and security, not for neural spike sorting.
- 6) *FPGA only (PG):* Same as *FPGA only (CG)* but with power gating during communication/security inactivity. While cardiac signal timing enables clock gating on the FPGA, power gating remains impractical due to the resulting millisecond-scale bitstream reconfiguration delays. Due to iCE40UP's SRAM-based architecture, an external flash memory is required for bitstream storage.

Fig. 3 summarizes the workloads and scenarios mentioned above. Table II describes the specifications of the assumed cardiac and neural IMDs.

F. Measurements

Current and workload execution time were measured using distinct methods for each device (see Table I).

1) CPU/MCU: For the EFM32TG11, live measurements were performed utilizing the advanced energy monitoring (AEM) feature of the SLSTK3301A starter kit and the Simplicity Studio IDE. The AEM infrastructure provides subsystem-level current monitoring of the EFM32TG11 device rather than total development-board consumption. Unused peripherals and on-board subsystems were disabled where possible during measurements. Additionally, one of the internal timers was employed to determine the run times of individual workloads.

TABLE III
FPGA RESOURCE UTILIZATION (iCE40UP5K)

Workload	LUTs	FFs	BRAM	DSP	I/O Buffers
Spike detector	249	124	0	5	14
Spike classifier	3184	1125	2	8	3
CSD	1034	619	0	2	11
AES-128	2026	833	0	0	13
SIMON/SPECK	305	214	0	0	12
PHOTON	276	188	0	0	12

2) FPGA: The iCE40 UltraPlus Breakout Board, which houses the iCE40UP5K FPGA, provides test points (TPs) connected to specific voltage lines for measuring current consumption via multimeters or oscilloscopes. Measurements focused on the FPGA-related supply rails (VCC) exposed through these test points rather than total board-level consumption. Unused peripherals and non-essential subsystems were assumed inactive during the experiments. The workload run times were calculated using RTL simulation. All FPGA implementations successfully met timing closure at 12 MHz under the standard iCE40UP5K synthesis and place-and-route flow; Table III summarizes the achieved resource utilization for each workload.

Because the MCU and FPGA measurements were performed on different development platforms, the reported results should be interpreted as processor-focused rather than complete board-level power comparisons. Whereas direct measurements were used for both MCU and FPGA implementations, eFPGA values were obtained from vendor-provided models as discussed in Section III-C.

G. Evaluation Metrics

A comparative analysis of CPU and FPGA implementations was conducted, examining **execution time**, **average power consumption** per workload iteration, **energy consumption** per iteration, and **power density**. Execution time, average power, and energy consumption are key metrics for evaluating real-time behavior and impact on device autonomy. Power density is crucial for determining the thermal profile and safety of IMDs implanted adjacent to human tissue, specifically regarding the risk of localized hot spot development. It should remain below the reported limit for clinical grade implants of approximately 62 mW/cm² [26].

Furthermore, a comprehensive assessment of **battery life** was conducted by comparing expected operational durations across diverse scenarios and CPU/FPGA configurations, taking into account workload demands (derived from physiological signal processing) and communication volume. Execution-time and power measurements were obtained directly from the evaluated workloads and platforms, whereas the parameters listed in Table II were incorporated into the higher-level battery-life and scenario analysis. Battery life was estimated from the aggregate energy consumption of active workload execution, idle/sleep intervals, wireless communication, and stimulation activity:

$$E_{\text{total}} = E_{\text{workload}} + E_{\text{idle}} + E_{\text{comm}} + E_{\text{stim}}.$$

III. RESULTS

A. Performance Analysis: CPU/MCU Vs. FPGA

Using the platform parameters outlined in Table II, Fig. 4 presents a comparison of iCE40UP and EFM32TG11 across the medical and security workloads selected in Section II-D. In the figure, the reported execution time per iteration corresponds only to the active workload execution interval and excludes idle/sleep periods between successive workload activations.

As anticipated, the iCE40UP demonstrates superior execution times compared to the EFM32TG11 across all workloads, attributed to its inherent parallelism. The sole exception is the hardware-accelerated AES-128 implementation within the EFM32TG11's cryptographic module, which outperforms the FPGA, as expected.

However, the average power consumption per workload iteration exhibits a more nuanced trend. For computationally intensive tasks, such as the MLP-based spike classifier and AES-128, increased parallelism in the FPGA results in a corresponding rise in power consumption.

Despite this power trade-off, the shorter execution times achieved by the FPGA significantly reduce overall energy consumption, leading to superior energy efficiency compared to the CPU-based implementations.

Power density trends closely align with power consumption, given the comparable die areas of the EFM32TG11 and iCE40UP. Notably, both devices remain below the prescribed power density limit of 62 mW/cm² (discussed in Section II-G).

B. Contextual Analysis: Battery Life & IMD Usage Scenarios

The results presented thus far suggest that FPGAs offer a performance advantage in IMD applications for many workloads. However, to comprehensively evaluate system-level implications, further investigation is required, specifically concerning: (1) the optimization of MCU-FPGA pairings, (2) the influence of realistic usage patterns in cardiac and neural settings, and (3) the sensitivity to data transmission volume. Fig. 5 illustrates the battery life characteristics of different IMD processing-device configurations under the four scenarios defined in Section II-E. The battery-life analysis accounts for workload execution energy, idle/sleep energy, communication energy, and stimulation energy using the assumptions summarized in Table II. Consequently, the observed battery-life improvements depend not only on gating opportunities, but also on workload execution efficiency and workload duty-cycle characteristics.

The analysis reveals that FPGA (CG/PG) configurations offer enhanced battery life compared to MCU-only configurations for neural workloads. However, this trend reverses for cardiac workloads, where MCUs exhibit superior performance. Specifically, FPGAs contribute to extended battery life in high-duty-cycle and highly-parallel neural applications, which offer limited gating opportunities. In these workloads, the primary benefit arises from reduced active execution time and lower energy consumption per workload iteration, rather than




Fig. 4. Performance comparison of EFM32TG11 (MCU) and ICE40UP (FPGA) across medical and security workloads. MCU(HW) denotes hardware-accelerated AES on the EFM32TG11. The FPGA achieves lower execution times for most workloads due to parallel execution. Although instantaneous power may increase in some cases, the reduced execution time often results in improved overall energy efficiency.




Fig. 5. Battery life of IMD processing-unit configurations across the four evaluated scenarios. Bars represent the nominal transmission case (5-MB daily for scenarios 1 and 3; weekly for scenarios 2 and 4), while whiskers indicate sensitivity to communication volume, spanning 10 MB (lower limit) and 1 MB (upper limit). Scenario definitions and processing-device configurations are provided in Section II-E. FPGA-based solutions improve battery life in neural workloads, whereas MCU-only solutions remain favorable in many cardiac scenarios due to limited FPGA power-gating opportunities.

aggressive power gating. Although FPGA power gating opportunities remain limited in neural workloads, the reduced execution time substantially decreases the active energy contribution to the overall battery budget. Conversely, MCUs demonstrate better performance in lower-duty-cycle cardiac workloads, primarily because the cardiac duty cycles are not low enough for FPGAs to employ power gating. However, when considering very-low duty cycle security workloads that *do* permit FPGA power gating, the performance of FPGAs and MCU-only implementations was comparable.

C. The Use of eFPGAs

In addition to neural workloads, the above analysis confirms the suitability of FPGAs for security workload implementation in modern IMDs. As a *complementary* investigation, we introduced an eFPGA-based security workload implementation for comparison. We used Menta's eFPGA Core IP [27], which was synthesized in TSMC 28 nm technology. Power consumption was estimated using Menta's Origami Designer and provided models, since physical current measurements were not possible due to the lack of a fabricated device. The eFPGA clock




Fig. 6. Performance comparison of eFPGA with EFM32TG11 (MCU) and iCE40UP (FPGA) across security workloads. The evaluated eFPGA implementation exhibits execution-time and power characteristics comparable to the standalone FPGA implementation.




Fig. 7. Impact of eFPGA on battery life. Because the evaluated eFPGA implementation demonstrates similar energy characteristics to the standalone FPGA, the resulting battery-life differences remain modest across the considered scenarios.

frequency and supply voltage were set to 12 MHz and 1.2 V, respectively, matching the iCE40UP.

Fig. 6 illustrates the performance comparison between the eFPGA/MCU hybrid (security-focused eFPGA), MCU-only, and MCU/iCE40UP solutions. While the eFPGA was expected to exhibit lower power consumption, its power usage is comparable to the iCE40UP, which is attributed to the considerable leakage current inherent in the 28 nm process.

Due to the similar power consumption and identical workload execution times observed for the eFPGA and iCE40UP, the resulting battery life impact is insignificant, as shown in Fig. 7.

IV. DISCUSSION

A. Regulatory Considerations

Our feasibility study would be incomplete without a thorough examination of the regulatory implications of integrating FPGAs into active IMDs. These devices fall under strict oversight by agencies such as the U.S. Food and Drug Administration (FDA) and the European Medicines Agency (EMA), which require

proof of safety, efficacy, and functional reliability. Any modification affecting a device's performance, whether hardware or software, typically mandates regulatory review.

IMDs such as pacemakers and neurostimulators are classified as *Class-III* devices under *MDR2017/745* due to their invasive, life-critical nature [28]. Certification involves rigorous evaluation by Notified Bodies in the EU [29] or by the FDA in the USA. FPGAs introduce flexibility via reconfigurable logic, enabling real-time optimizations post-implantation, yet regulatory authorities treat any functional change as a potential trigger for re-evaluation [30]. Key determinants include: (1) the extent of functional change, (2) its impact on patient safety, treatment efficacy, and device reliability, and (3) the validation and verification processes assessing FPGA reconfiguration.

1) In-Vivo Reconfiguration Implications: The regulatory response to in vivo FPGA reconfiguration depends on the modification's scope. Major changes that alter core functionality or introduce risk require review, whereas pre-approved updates within defined safety boundaries may proceed without full recertification [31]. Notably, **regulators treat FPGA configurations analogously to software**, meaning that their reconfigurability does not automatically trigger additional certification hurdles [32], [33]; emphasis lies on verification of function rather than implementation details. FPGAs are generally regarded as “black boxes” unless they perform safety-critical operations [32]. Factors influencing approval include:

- Whether FPGA modifications remain within an approved operational envelope.
- The device's ability to maintain consistent, validated performance across configurations.
- The regulatory classification of the IMD and degree of control over in-field updates.

2) Recertification Triggers and Regulatory Pathways: Recertification requirements depend on the update magnitude:

Minor updates (e.g., bug fixes or algorithm refinements) may require documentation, *notification* or limited review depending on the applicable regulatory pathway [34]. **Major updates** that affect safety or efficacy may trigger *re-evaluation* within existing certification cycles, typically every five years under MDR2017/745, Annex VII, Section 4.11 [28].

Regulatory frameworks offer established pathways for managing such updates: The FDA allows Premarket Approval (PMA) supplements or Premarket Notification (510(k)) submissions for significant changes [35], while preapproved update protocols may reduce or avoid the need for complete recertification if FPGA changes remain within verified design boundaries [31]. Risk-based assessments under ISO 14971 guide whether reconfiguration warrants additional filings [36].

3) Applicable Standards and Verdict: Finally, *expert commentary and industry sources* suggest that FPGA updates follow the same rules as software updates, without additional obligations compared to MCU-based systems [32], [34], [37], [38]. This is confirmed in the IEC 62304 standard, governing software life-cycle processes, which implicitly covers FPGA code as a form of complex design unless the design qualifies as pure (e.g., easy-to-test) hardware under IEC 60601-1 [33], [39]. Its 2015 update formalized FPGAs within the software life-cycle framework, closing any potential regulatory loopholes.

B. Application Suitability: Cardiac Vs. Neural IMDs

The analysis suggests that FPGAs may not be optimal for cardiac-IMD medical functions, primarily because the duty cycles enforced by the physiological deadlines are not low enough to allow FPGA power gating, unlike MCU implementations. However, FPGAs demonstrate significant potential for neurostimulators. This is particularly advantageous given the parallel and computationally intensive algorithms employed in neurostimulation, which must process rapidly changing and often noisy neural signals.

C. Extended Battery Life in Neurostimulators

By leveraging the power efficiency of FPGAs for neural workloads, it is feasible to design neurostimulators that are either non-rechargeable or require significantly less frequent recharging. This extended operational lifespan is crucial for patient comfort and reduced maintenance.

D. Enhanced Security and Upgradability

Even when the medical algorithm is implemented on an MCU, FPGAs offer a compelling solution to implement robust security schemes. This is especially relevant because security algorithms often require frequent updates to address evolving threats. Accelerating these security workloads, particularly with power gating, proves to be highly effective.

Besides, emerging ultra-low-power AI/ML accelerators may offer improved efficiency for specific inference tasks, compared to FPGAs. However, such architectures are often optimized for narrow workload classes and may not provide the flexibility

TABLE IV
SUMMARY OF RESULTS

Workload Type	Duty Cycling	Deepest Sleep Mode		FPGA Feasibility
		MCU	FPGA	
Cardiac	Medium	PG	CG	Less Feasible
Neuro	High	N/A	N/A	Feasible
Security	Low	PG	PG	Feasible

required to support the broader mix of medical, security, and adaptive workloads targeted in this study.

E. Limitations of Power Gating in Medical Workloads

The inherent constraints imposed by the real-time requirements of physiological signals, specifically in medical workloads, preclude the effective use of power gating for algorithm acceleration. The timing demands of these signals do not allow sufficient time for SRAM-based FPGA reconfiguration from external flash, rendering power gating impractical for neural and many cardiac workloads.

F. Power-Density Compliance

The power density observed in our FPGA-based IMD implementations remains well below the clinically prescribed upper limit of approximately 62 mW/cm² [26]. This ensures compatibility with stringent safety standards for implantable medical devices.

Table IV summarizes this analysis, which underscores the suitability of FPGAs for compute-intensive IMD workloads characterized by high duty cycles. Furthermore, FPGAs demonstrate optimal performance in scenarios with extremely sparse usage, particularly security workloads, due to the applicability of power gating. Medium duty cycle applications, such as those in cardiac IMDs, could theoretically benefit from power gating; however, this is currently impractical due to the high reconfiguration latency of SRAM-based FPGAs. Future IMD designs incorporating flash-based FPGAs may overcome these limitations, enabling reconfiguration with significantly lower energy and performance overhead.

G. Study Limitations

The presented operating scenarios in this work should be interpreted as representative and exploratory rather than clinically predictive. Accordingly, the reported results should be viewed as directional for the evaluated design space, rather than universally predictive across all implant configurations, clinical use cases, or future low-power processing architectures. The analysis is intended to expose system-level architectural trends associated with alternative processing cores under consistent assumptions, rather than to provide a complete clinical deployment model for a specific implant product.

The presented conclusions are also derived from a limited set of evaluated platforms, namely a specific MCU, FPGA family, and eFPGA implementation operating under representative assumptions and workload scenarios. Consequently, the quantitative results should not be interpreted as universally generalizable across all future low-power reconfigurable architectures, process

nodes, implant SoCs, or emerging heterogeneous accelerator platforms. Instead, they should be understood as comparative evidence of the architectural trade-offs observed within the evaluated design space.

Similarly, the battery-life estimates presented in this work should be interpreted as first-order, comparative projections rather than clinically predictive estimates. The model accounts for workload execution, idle/sleep behavior, wireless communication, and stimulation activity, but does not explicitly model advanced battery aspects, such as aging, temperature effects, long-term leakage evolution, and so on. These factors may affect absolute device lifetime and, in some regimes, stimulation energy may dominate computational energy. However, because stimulation and communication assumptions are held fixed across the processing configurations within each scenario, the reported comparisons primarily capture the relative contribution of the processing subsystem. The resulting trends should therefore be interpreted as directional architectural trade-offs under representative assumptions, rather than full electrochemical or clinical battery-life predictions.

V. CONCLUSION

This work has demonstrated that FPGAs can serve as effective accelerators for IMDs, particularly for neural and security-intensive workloads. Our analysis indicates that while FPGAs may not be optimal for cardiac IMDs, mainly because the duty cycles dictated by physiological deadlines are insufficiently low to permit effective power gating, they exhibit significant potential for neurostimulation systems. This advantage arises from their ability to efficiently execute parallel, computation-intensive algorithms required to process rapidly fluctuating and often noisy neural signals. As a result, FPGA-based implementations can enhance device longevity by extending battery life and reducing recharge frequency, thereby improving patient comfort and minimizing maintenance requirements. Security remains a critical concern for IMDs, and FPGAs offer a robust platform for cryptographic acceleration. The results show that FPGAs can efficiently offload security tasks and maintain adaptability to evolving security threats through their reconfigurable nature. Furthermore, they comply with strict power-density and safety constraints essential for IMD applications. Future work will focus on exploring partial reconfiguration, flash-based FPGA technologies, hybrid FPGA-GPU architectures, and the deployment of more advanced AI models, including CNNs, spiking neural networks, adaptive online-learning pipelines, and larger neural-processing workloads, to further optimize performance and energy efficiency in next-generation IMDs.

REFERENCES

- [1] L. F. Herbozo Contreras et al., "Neuromorphic neuromodulation: Towards the next generation of closed-loop neurostimulation," *PNAS Nexus*, vol. 3, no. 11, 2024, Art. no. pgae488.
- [2] T. Harland, T. Elliott, I. Telkes, and J. G. Pilitsis, "Machine learning in pain neuromodulation," in *Computat. Neurosurgery*, pp. 499–512, 2024.
- [3] D. De Ridder, M. A. Siddiqi, J. Dauwels, W. A. Serdijn, and C. Strydis, "NeuroDots: From single-target to brain-network modulation: Why and what is needed," *Neuromodulation: Technol. Neural Interface*, vol. 27, no. 4, pp. 711–729, 2024.
- [4] Y. Li et al., "Brain-machine interactive neuromodulation research tool with edge AI computing," *Heliyon*, vol. 10, no. 12, 2024.
- [5] D. M. Corva et al., "SmartStim: An artificial intelligence enabled deep brain stimulation device," *IEEE Trans. Med. Robot. Bionics*, vol. 6, no. 2, pp. 674–684, May 2024.
- [6] M. A. Siddiqi, G. Hahn, S. Hamdioui, W. A. Serdijn, and C. Strydis, "Improving the Security of the IEEE 802.15. 6 Standard for Medical BANs," *IEEE Access*, vol. 10, pp. 62953–62975, 2022.
- [7] M. A. Siddiqi, C. Doerr, and C. Strydis, "IMDfence: Architecting a secure protocol for implantable medical devices," *IEEE Access*, vol. 8, pp. 147948–147964, 2020.
- [8] M. Sawan, S. Robin, B. Provost, Y. Eid, and K. Arabi, "A wireless implantable electrical stimulator based on two FPGAs," in *Proc. Third Int. Conf. Electronics, Circuits, Syst.*, vol. 2, 1996, pp. 1092–1095.
- [9] T.-t. Hoang, J.-p. Son, Y.-r. Kang, C.-r. Kim, H.-y. Chung, and S.-w. Kim, "A low complexity, low power, programmable QRS detector based on wavelet transform for implantable pacemaker IC," in *Proc. IEEE Int. SOC Conf.*, 2006, pp. 160–163.
- [10] Silicon Laboratories Inc., *EFM32 Tiny Gecko 11 Family Reference Manual*, 2025.
- [11] M. I. Khan and B. da Silva, "Harnessing FPGA technology for energy-efficient wearable medical devices," *Electronics*, vol. 13, no. 20, 2024, Art. no. 4094. [Online]. Available: <https://www.mdpi.com/2079-9292/13/20/4094>
- [12] M. Vaithianathan, M. Patil, S. F. Ng, and S. Udkar, "Energy-efficient FPGA design for wearable and implantable devices," *ESP Int. J. Advancements Sci. Technol.*, vol. 2, no. 2, pp. 37–51, 2024. [Online]. Available: <https://www.espjournals.org/IJAST/ijast-v2i2p107>
- [13] M. B. Altman, W. Wan, A. S. Hosseini, S. A. Nowdeh, and M. Alizadeh, "Machine learning algorithms for FPGA implementation in biomedical engineering applications: A review," *Heliyon*, vol. 10, no. 4, 2024, Art. no. e26652, DOI: [10.1016/j.heliyon.2024.e26652](https://doi.org/10.1016/j.heliyon.2024.e26652).
- [14] Lattice Semiconductor Corporation, *iCE40 UltraPlus™ Family*, 2017.
- [15] Microsemi, *ZL70103 Medical Implantable RF Transceiver*, 2015.
- [16] Macronix International Co., Ltd., *MX25V1006F 1M-BIT [x 1/x 2] CMOS SERIAL NOR FLASH - Datasheet*, 2023.
- [17] M. A. Siddiqi et al., "A lightweight architecture for real-time neuronal-spike classification," in *Proc. 21st ACM Int. Conf. Comput. Front.*, 2024, pp. 32–40.
- [18] Y. Chen, M. Dhamala, A. Bollimunta, C. E. Schroeder, and M. Ding, "Current source density analysis of ongoing neural activity: Theory and application," in *Electrophysiological Recording Techn.*, Berlin, Germany: Springer 2011, pp. 27–40.
- [19] Y. Yang and A. J. Mason, "Hardware efficient automatic thresholding for neo-based neural spike detection," *IEEE Trans. Biomed. Eng.*, vol. 64, no. 4, pp. 826–833, Apr. 2017.
- [20] PLC Medtronic, *Activa™ PC Multi-Program Neurostimulator - Device Manual*, 2020.
- [21] M. Deterre, "Toward an energy harvester for leadless pacemakers," Ph.D. dissertation, Université Paris Sud - Paris XI, Orsay, France, Jul. 2013. [Online]. Available: <https://tel.archives-ouvertes.fr/tel-00868838>
- [22] J. Nielsen, A. Pedersen, P. Mortensen, and H. Andersen, "Programming a fixed long atrioventricular delay is not effective in preventing ventricular pacing in patients with sick sinus syndrome," *EP Europace*, vol. 1, no. 2, pp. 113–120, 1999.
- [23] F. Simon, J. P. Martinez, P. Laguna, B. van Grinsven, C. Rutten, and R. Houben, "Impact of sampling rate reduction on automatic ECG delineation," in *Proc. 29th Annu. Int. Conf. IEEE Eng. Med. Biol. Soc.*, 2007, pp. 2587–2590.
- [24] NeuroPace, *RNS® System User Manual*, 2019.
- [25] M. Balish, P. S. Albert, and W. H. Theodore, "Seizure frequency in intractable partial epilepsy: A statistical analysis," *Epilepsia*, vol. 32, no. 5, pp. 642–649, 1991.
- [26] F. Zhang, M. Aghagolzadeh, and K. G. Oweiss, "A fully implantable, programmable and multimodal neuroprocessor for wireless, cortically controlled brain-machine interface applications," *J. Signal Process. Syst.*, vol. 69, pp. 351–361, 2012, Art. no. 5293973. [Online]. Available: <https://api.semanticscholar.org/CorpusID>
- [27] M. SAS, *eFPGA Core IP - Product Brief*, 2016.
- [28] European Commission, "Guidance on qualification and classification of software in regulation (EU) 2017/745-MDR and regulation (EU) 2017/746 - IVDR," 2019. [Online]. Available: https://health.ec.europa.eu/system/files/2020-09/md_mdcg_2019_11_guidance_en_0.pdf
- [29] European Commission, "Notified bodies," https://single-market-economy.ec.europa.eu/single-market/goods/building-blocks/notified-bodies_en, 2023. Accessed: Mar. 29, 2025.

- [30] *Int. Med. Device Regulators Forum (IMDRF)*, "Software as a medical device (SaMD): Possible framework for risk categorization and corresponding considerations," 2014. [Online]. Available: www.imdrf.org/documents/software-medical-device-possible-framework-risk-categorization-and-corresponding-considerations
- [31] U. S. Food and Drug Administration, *General Princ. of Soft. Validation: Final Guid. for Industry and FDA Staff*, 3rd ed. Center for Devices and Radiological Health, U.S. FDA, Jan. 2002, version 2.0 (original: 1997), updated in 2021 for editorial changes. [Online]. Available: <https://www.fda.gov/media/73141/download>
- [32] Blue Pearl Software, "RTL development and testing for medical devices," 2017. [Online]. Available: <http://www.bluepearlsoftware.com/files/MedicalDevicesPaperFinal.pdf>
- [33] Medical device software—Software life cycle processes, International Organization for Standardization Std., IEC 62304:2006 Amd 1:2015, Apr. 2015. [Online]. Available: <https://www.iso.org/standard/64686.html>
- [34] P. J. Tanzillo, "Controlling risk in medical devices," 2009. Accessed: Apr. 4, 2021. [Online]. Available: <https://www.machinedesign.com/archive/article/21829447/controlling-risk-in-medical-devices>
- [35] U. S. Food and Drug Administration. (2025) *Premarket Approval (PMA)*, Center for devices and radiological health (CDRH). [Online]. Available: <https://www.fda.gov/medical-devices/premarket-submissions-selecting-and-preparing-correct-submission/premarket-approval-pma>
- [36] Medical devices — Application of risk management to medical devices International Organization for Standardization Std., ISO 14971:, 2019 edition: 3, Geneva, Switzerland, 2019. [Online]. Available: <https://www.iso.org/standard/72704.html>
- [37] M. E. Azzouzi, "Classifying medical devices according to MDR2017/745," 2021. Accessed: Jun. 21, 2021. [Online]. Available: <https://www.easymedicaldevice.com>
- [38] J. Gable, "Does FPGA code count as software," 2020, Accessed: Jun. 4, 2021. [Online]. Available: <https://jeffgable.com/daily/does-fpga-code-count-as-software/>
- [39] International Electrotechnical Commission, Medical Electrical Equipment—Part 1: General Requirements for Basic Safety and Essential Performance, Standard, International Electrotechnical Commission Std. IEC 60601-1, Geneva, Switzerland, Aug. 2015. [Online]. Available: <https://webstore.iec.ch/publication/21967>