Energy-Efficient Frequency Selection Method for Bio-Signal Acquisition in AI/ML Wearables

Hossein Taji[1], José Miranda[1], Miguel Peón-Quirós[2], David Atienza[1,2]
{hossein.taji,jose.mirandacalero,miguel.peon,david.atienza}@epfl.ch
[1]Embedded Systems Laboratory (ESL), École Polytechnique Fédérale de Lausanne (EPFL), Switzerland
[2]EcoCloud, École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

ABSTRACT

In wearable sensors, energy efficiency is crucial, particularly during phases where devices are not processing, but rather acquiring biosignals for subsequent analysis. This study focuses on improving the power consumption of wearables during these acquisition phases, a critical but often overlooked aspect that substantially affects overall device energy consumption, especially in low-duty-cycle applications. Our approach optimizes power consumption by leveraging application-specific requirements (e.g., required signal profile), platform characteristics (e.g., transition-time overhead for the clock generators and power-gating capabilities), and analog biosignal front-end specifications (e.g., ADC buffer sizes). We refine the strategy for switching between low-power idle and active states for the storage of acquired data, introducing a novel method to select optimal frequencies for these states. Based on several case studies on an ultra-low power platform and different biomedical applications, our optimization methodology achieves substantial energy savings. For example, in a 12-lead heartbeat classification task, our method reduces total energy consumption by up to 58 % compared to state-of-the-art methods. This research provides a theoretical basis for frequency optimization and practical insights, including characterizing the platform's power and overheads for optimization purposes. Our findings significantly improve energy efficiency during the acquisition phase of wearable devices, thus extending their operational lifespan.

CCS CONCEPTS

• **Computer systems organization** → **Embedded systems**; • **Hardware** → *Signal processing systems* .

KEYWORDS

AI/ML Wearables, Energy Efficiency, Bio-signal Acquisition, Frequency Optimization, Biomedical Signal Processing

ACM Reference Format:

Hossein Taji[1], José Miranda[1], Miguel Peón-Quirós[2], David Atienza[1,2]. 2024. Energy-Efficient Frequency Selection Method for Bio-Signal Acquisition in AI/ML Wearables . In *Proceedings of the ACM/IEEE International Symposium on Low Power Electronics and Design (ISLPED '24), August 5–7, 2024, Newport Beach, CA, USA.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3665314.3670815

ISLPED '24, August 5–7, 2024, Newport Beach, CA, USA

ACM ISBN 979-8-4007-0688-2/24/08…$15.00
https://doi.org/10.1145/3665314.3670815

1 INTRODUCTION

The increasing healthcare costs and the essential need for preventive measures are driving the development of wearables for monitoring and processing long-term and real-time biosignals [15]. To ensure user privacy and extend battery life, it is vital to perform the entire processing chain directly on the wearable device.

RISC-V-based Ultra-Low-Power (ULP) platforms such as [9, 10, 16, 18], using an open-source ISA for flexibility, rapid innovation and energy efficient design, offer the necessary computational power and memory capabilities for ML/AI applications.

Within this context, a portion of an application's lifecycle is not dedicated to processing but to the acquisition of bio-signals for subsequent analysis. This acquisition phase is particularly critical in low-duty-cycle applications, where it accounts for a substantial part of the energy consumption. For example, [1] demonstrates an extremely low duty-cycle scenario, with bio-signals collected for 1 s and processing that lasts less than 1 ms. Although the importance of energy efficiency during the processing phase is well recognized, particularly in the utilization of heterogeneous platforms [6], the energy consumed during the acquisition phase often remains overlooked and unreported [1, 7, 12]. This gap highlights a critical area for energy optimization that has not yet been fully explored and addressed in wearable healthcare technologies.

A balanced approach is necessary to enhance energy efficiency during the acquisition phase, incorporating power gating and frequency reduction strategies. Power gating can significantly reduce both static and dynamic power consumption by shutting off non-essential modules, whereas frequency reduction aims at minimizing dynamic power. Two prevalent methods emerge in the literature: maintaining only essential modules for sample reception with the remainder of the system powered down [14, 20, 21], and switching between an extremely low-power idle state, and an active state upon sample reception [1, 5]. The latter requires careful consideration of optimal frequencies to minimize energy consumption, considering platform characteristics such as static and dynamic power in each state, transition time overheads from FLL and power gating, buffer size of Analog Front Ends (AFE), and an application's required number of channels and signal frequencies.

Our contributions to enhancing wearable device energy efficiency during biosignal acquisition phases include 1) a novel methodology for optimal frequency selection in idle and active states, tailored to both uniform and dynamic frequency adjustments under varied overhead conditions; 2) a comprehensive system model for power optimization based on target frequencies; 3) in-depth characterization of a case study platform, including FLL overhead analysis and modeling, and power assessment; 4) Practical validation via different biomedical applications, demonstrating up to 58 % energy savings compared to existing methods.

2 RELATED WORK

There are three main approaches when dealing with the acquisition phase: 1) All-On, where there is no distinct policy differentiating signal acquisition from processing; 2) Continuous Streaming (CS), which employs a continuous low-power mode for sample storage; and 3) Burst, which alternates between an active state for storing samples and a very low-power idle state between samples.

Different works leverage the All-On strategy. For instance, [12], [11], [13], and [7] use the Mr. Wolf platform [18] for Brain-Computer Interface (BCI), 4-lead EEG-based epilepsy monitoring and 8-lead and 16-lead EMG-based gesture recognition, maintaining the same power profile during acquisition as in processing. Similarly, [8] uses GAP9 [10] for EEG- and PPG-based BCI. However, these works do not adopt specialized policies for the acquisition phase, resulting in similar power usage for both the acquisition and processing phases.

For the second approach, [20] examines 8-lead sEMG-based hand movement classification on GAP8 [9], adopting a low-power mode by clock-gating the 8-core cluster and keeping the MCU active for signal storage. [21] and [14] use similar strategies on Mr. Wolf for 1-lead ECG-based QRS complex detection and 3-lead EEG-based drowsiness detection, with [21] also sleeping the CPU during acquisition, a method not specified in [14]. However, maintaining a constant mode during acquisition prevents these studies from exploring potentially more energy-efficient strategies, such as power-gating non-essential components like DMA and SPI interfaces, during intervals between sample receptions.

For the burst approach, [1] and [5] explore EEG-based epileptic seizure detection and EMG-based gesture recognition on Mr. Wolf, respectively. Their approach alternates between low-power idle and active states in response to incoming samples. During idle periods, most of the system, including the MCU and cluster, enters a low-power mode, significantly cutting energy use. Upon receiving samples, it transitions to an active state, activating the MCU for data storage. Although these studies leverage reduced power by toggling between two states, they adopt a straightforward frequency selection strategy: opting for the lowest available frequency on the platform during idle periods and maximizing frequency within the minimal voltage range for active periods [1, 5].

It is important to note that AFEs often include built-in small memories for temporary sample storage. For example, [2] and [3] are equipped with FIFO buffers capable of holding 128 and 256 samples of 8 bit and 32 bit, respectively. Although these buffers, which introduce minimal power overhead [3], can significantly reduce the frequency of data transmission, their advantages remain underexploited in the discussed wearable technologies.

We propose a novel optimization method that minimizes the burst-mode power consumption by selecting optimal frequencies for idle and active states. To our knowledge, this is the first proposition of such a methodology. We demonstrate substantial energy savings over conventional approaches in different case study applications.

3 METHODOLOGY FOR OPTIMIZATION OF ACQUISITION POWER

Our methodology aims to reduce energy consumption in wearables by optimizing frequency selection for the acquisition phase, which can be challenging due to the wide range of frequencies supported at each voltage level. This platform-independent procedure, applicable to any setup, begins by modeling various factors influencing optimal frequency selection, as shown in Figure 1. Note that this step is done once per system. Then, we outline steps for extracting an accurate wearable platform characterization for determining values like timing overheads and power profiles. The final phase uses these parameters for optimization considering two scenarios: dynamically changing frequency or maintaining a consistent frequency in burst mode, where the optimal choice may vary based on application. This process is assessed under different settings, for example, to see whether the time required by the FLL to switch from one frequency to another depends on the distance between both frequencies. Since CS mode's optimal frequency is fixed by the application bandwidth, we focus our analysis on optimizing the frequencies for the active (F_a) and idle (F_i) states of the burst mode.

This methodology is especially useful for dynamic applications, where the setup may change during runtime, such as the number of used signal channels or algorithm complexity. For such applications, the methodology can be implemented to calculate optimal configurations online, or pre-calculate optimal configurations for different scenarios offline during design time.

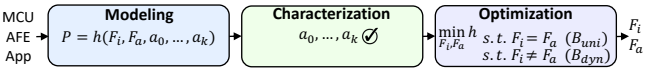

Figure 1: Methodology flow for optimizing acquisition power.

3.1 System Modeling

The initial step involves system modeling to accurately represent the wearable device architecture, typically comprising an MCU connected to an AFE for biosignal reception, as depicted in Figure 2. AFEs often support multiple channels for simultaneous bio-signal capture and may include a FIFO buffer for temporary data storage before transmission. The key to modeling the power consumption of the MCU during acquisition is identifying essential parameters. The application's bandwidth (BW) is dictated by the number and sampling frequency of required biosignals. The burst frequency (F_b), or how often the AFE signals the MCU to obtain data, is inversely related to the buffer size (N_s). Thus, assuming optimal buffer utilization, we define the burst frequency as $F_b = \frac{BW}{N_s}$.

The power in burst mode, P_B, represents the average power over the interval between bursts ($T_b = 1/F_b$), accounting for both active (t_a) and idle (t_i) periods:

$$P_B = \frac{P_a t_a + P_i t_i}{t_a + t_i} = F_b(P_a t_a + P_i t_i) = F_b t_a (P_a - P_i) + P_i. \quad (1)$$

where P_a and P_i indicate power in active and idle states.

As we set our optimization objective based on the system's frequencies in idle, F_i, and active, F_a states, it is imperative to define the power consumption for each state, P_i and P_a, as a function of

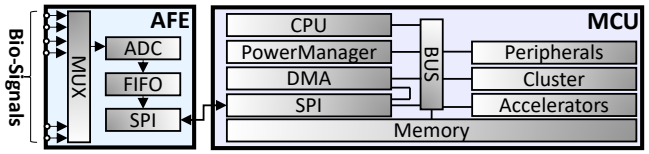

Figure 2: A common wearable device architecture.

their respective frequencies. CMOS circuit power consumption is typically divided into static (P_S) and dynamic (P_D) components, where $P = P_S + P_D$ [17]. Unlike P_S, which is consistent across frequencies, P_D varies linearly with system frequency (F), influenced by the activity factor (α), load capacitance (C), and the voltage squared (V^2), hence $P_D = \alpha C V^2 F$ [17]. By introducing the dynamic power coefficient, $\eta = \alpha C V^2$, the power consumption at any given frequency for idle and active states can be expressed as follows:

$$P_i = P_{i_s} + \eta_i F_i, \quad P_a = P_{a_s} + \eta_a F_a. \tag{2}$$

To account for the timing factors affecting idle and active durations, we examine the overheads arising from power gating and reactivating modules (t_p), and frequency adjustments via FLL (t_f). In the acquisition phase, the main task is transferring samples from the AFE buffer to platform memory, typically facilitated by DMA. This transfer time (t_s) is directly proportional to the volume of data to be transferred, represented by the number of samples (N_s), and inversely proportional to the system's operational frequency during active state (F_a). Thus, we represent t_s with the equation $t_s = \beta \frac{N_s}{F_a}$, where β acts as a proportionality constant, capturing the operational overheads and transfer efficiencies.

Considering these dynamics, the active period duration (t_a) integrates these timing components as follows:

$$t_a = \beta \frac{N_s}{F_a} + t_p + t_f. \tag{3}$$

Merging this with the power components from Equations 2 and 1, we derive the formula for the power consumption in burst mode:

$$P_{\mathrm{B}} = F_b(\beta \frac{N_s}{F_a} + t_p + t_f)((P_{a_s} - P_{i_s}) + \eta_a F_a - \eta_i F_i) + (P_{i_s} + \eta_i F_i). \tag{4}$$

This equation frames our optimization problem, aiming to find F_a and F_i that minimize overall power consumption in burst mode.

Our optimization primarily targets the burst mode, given its intricate dynamics and power-saving potential. However, it should be noted that the power consumption of the CS mode follows a straightforward formula, since the system is always completely active: $P_{\mathrm{cs}} = P_{\mathrm{cs}_s} + \eta_{\mathrm{cs}} F_{\mathrm{cs}}$.

3.2 System Characterization

The system characterization involves extracting the timing and power parameters of the system required by our formulation, particularly timing overheads for FLL reconfiguration (t_f) and module power transitions (t_p). Accurate modeling of t_f, crucial due to its potential variation with the frequencies F_a and F_i, requires measuring transition times across a broad frequency range, encompassing both low and high spectrums. Analysis of these measurements can inform the development of a predictive model through regression or other fitting techniques. Additionally, the transfer efficiency constant, β, is determinable by measuring sample transfers from AFE to MCU at different frequencies.

Choosing which system modules to keep active and which ones to power down is also important. In CS mode, only those modules essential to implement continuous sample acquisition need to be active, with others power-gated. In contrast, during burst mode's idle state, all modules except those essential for waking up and maintaining data integrity are power gated to achieve minimal power; in burst mode, the CPU and most of the rest of the system modules need to be activated only during the active state to respond to the interrupts generated by the AFE.

Accurate optimization requires detailed energy profiles, including static and dynamic components, for each custom-defined acquisition mode. It requires individual power consumption measurements of each system module. However, directly measuring module-specific power in a circuit is often unfeasible. An effective alternative involves utilizing post-place-and-route (post-PNR) analysis. This technique uses the netlist to derive the relative power contributions of each module, allowing an estimation of individual module power based on the total system power measured on the actual platform. For separating the static (P_S) and dynamic power coefficient (η), we measure the power across frequencies at a constant voltage for a specific application, then fit a linear model, according to Equation 2, to these measurements. In scenarios lacking netlist access, estimations must rely on best-guess approximations or published data for similar components in analogous technologies.

3.3 System Optimization

Our objective is to minimize power consumption by optimizing the system frequencies (F_a and F_i) in Equation 4 under two scenarios of uniform or dynamic frequencies and various settings, especially considering t_f's impact, which can vary between platforms.

Uniform Frequency Optimization: If the system uses always the same frequency, $F = F_i = F_a$, the FLL configuration time is zero ($t_f = 0$). This modifies Equation 4 as follows:

$$P = \frac{a_0}{F} + a_1 F + a_2, \tag{5}$$
$$\text{where } a_0 = F_b \beta N_s (P_{a_s} - P_{i_s}), \quad a_1 = \eta_i + F_b(t_p)(\eta_a - \eta_i),$$
$$a_2 = F_b \beta N_s (\eta_a - \eta_i) + P_{i_s} + F_b(t_p)(P_{a_s} - P_{i_s})$$

Optimizing for minimum power consumption leads us to find the optimal frequency, F^*, by setting the derivative of P with respect to F to zero, $P' = a_1 F^2 - a_0 = 0$, which gives:

$$F^* = \sqrt{\frac{a_0}{a_1}} = \sqrt{\frac{F_b \beta N_s (P_{a_s} - P_{i_s})}{\eta_i + F_b t_p (\eta_a - \eta_i)}}. \tag{6}$$

Dynamic Frequency Optimization, General Case: In the case where frequencies dynamically switch between idle and active states, and the transition time t_f depends on these frequencies ($t_f = g(F_a, F_i)$), the power consumption P becomes a complex function of multiple terms involving F_a and F_i. For instance, modeling t_f as a linear function of the frequency difference introduces terms such as F_a^2 and F_i^2 and interactions between F_a and F_i, resulting in $P = h_1\left(F_a, F_i, \frac{1}{F_a}, F_a^2, F_i^2, F_a F_i, \frac{F_i}{F_a}\right)$. Due to this complexity, optimization requires numerical methods.

Even simplifying the scenario by fixing F_i at the platform's lowest frequency ($F_i = F_l$) leads to a reduced complexity in P, expressed as $P = h_2\left(F_a, \frac{1}{F_a}, F_a^2\right)$. However, this simplification still necessitates numerical solutions due to the cubic equation from the derivative with respect to F_a. This highlights that, in this case, an analytical solution is not feasible and the designer will need to use numerical methods on Equation 4 with the parameters of the concrete system to reach optimized frequencies.

Dynamic Frequency Optimization, Constant Overhead: For platforms where the timing overhead is constant $t_a + t_f = C$, as F_i solely influences P_i in Equation 1, we should put F_i to the lowest frequency our platform supports, $F_i = F_l$, to minimize P. Given this, Equation 4 is modified as follows:

$$P = \frac{a_3}{F_a} + a_4 F_a + a_5, \tag{7}$$
$$\text{where } P_{i,l} = P_{i_s} + \eta_i F_l,$$
$$a_3 = F_b \beta N_s (P_{a_s} - P_{i,l}), \quad a_4 = F_b (t_p + t_f)\eta_a,$$
$$a_5 = P_{i,l} + F_b \beta N_s \eta_a + F_b (t_p + t_f)(P_{a_s} - P_{i,l}).$$

Optimization targets finding the optimal active frequency, F_a^*, by equating the derivative of P with respect to F_a to zero, resulting in:

$$F_a^* = \sqrt{\frac{a_3}{a_4}} = \sqrt{\frac{\beta N_s (P_{a_s} - P_{i,l})}{(t_p + t_f)\eta_a}}. \tag{8}$$

Dynamic Frequency Optimization, No Overhead: In an ideal case without transition and configuration overheads ($t_a + t_f \approx 0$), the optimal frequency, according to Equation 8, trends towards the platform's maximum capability. Thus, for maximum energy efficiency, setting F_a to the highest supported frequency (F_h) at the platform's lowest voltage is recommended.

In summary, optimal frequency selection depends on whether frequencies between states are uniform ($F_a = F_i$) or vary ($F_i \neq F_a$). Uniform frequency optimization is guided by Equation 6. For dynamic frequencies, general scenarios with frequency-dependent overheads require numerical methods like Gradient Descent for power minimization. With fixed overheads, Equation 8 determines F_a with F_i set to the lowest frequency (F_l). Ideally, with minimal overheads ($t_a + t_f = 0$), frequencies are set to their operational extremes ($F_a = F_h$ and $F_i = F_l$) for maximum efficiency.

4 EXPERIMENTAL SETUP

4.1 Case Study Applications

This work focuses on two applications as case studies, chosen for their representation of scenarios with low-duty cycles where acquisition power significantly impacts overall energy efficiency. The first application is heartbeat classification (HBC), using electrocardiogram (ECG) signals for real-time detection of abnormal heart patterns indicative of common heart diseases. This application is explored in three configurations: 1-lead, 3-lead, and 12-lead ECG, each sampled at 256 Hz, as detailed in [6]. HBC also showcases a dynamic application, where the number of leads may change during runtime, which highlights the utility of our methodology for dynamic applications. The second application, real-time cognitive workload monitoring (CWM), uses electroencephalogram (EEG) signals to monitor cognitive states on wearable devices, showcasing its utility through a four-lead EEG setup sampled at 256 Hz [19]. The selection of HBC and CWM as case studies highlights examples of applications where efficient power management during acquisition is crucial because of their inherently low-duty cycles.

4.2 Case Study Platform: HEEPocrates

Our experimental setup utilizes the HEEPocrates platform [16], although our platform-independent methodology can be applied to any hardware setup. HEEPocrates is a system evolving from the X-HEEP architecture optimized for ultra-low-power edge computing in healthcare. This RISC-V-based platform integrates a host processor with specialized accelerators for a diverse computational ecosystem. HEEPocrates employs an advanced power management strategy, which includes clock-gating, power-gating, and RAM retention for detailed power control. However, certain chip domains, such as the always-on peripheral, cannot be power-gated. The platform design supports dynamic voltage and frequency scaling through an FLL, facilitating flexible frequency adjustments for acquisition and processing.

In addition, this section covers two key characterizations of the HEEPocrates platform: FLL timing analysis and modeling, and detailing the energy consumption profiles for each acquisition mode.

Table 1: Static power consumption (P_s) and dynamic power coefficients (η) of HEEPocrates for different acquisition modes at 0.8 V.

Mode	P_s (µW)	η(µW/MHz)
Continuous Sampling (CS)	231.66	10.49
Burst Mode (Idle)	117.77	2.24
Burst Mode (Active)	349.20	22.83

4.2.1 Energy Profiles of HEEPocrates. For power profile assessment in each acquisition mode on the HEEPocrates platform [16], it is crucial to first determine which modules are used per mode. In CS mode, for energy-efficient data streaming via DMA, non-essential modules like the CPU, debug unit, PLIC, GPIO, timers, I2C, UART, and accelerators are power-gated. Conversely, burst mode's idle state retains only essential modules: bus, memory, FLL, Boot ROM, GPIO, SoC controller, fast interrupt controller, to ensure system wake-up and prevent data loss. During the active state of burst mode, the CPU and most modules are reactivated, except for accelerators.

To assess the power profiles of these modes, we combined empirical measurements across the voltage range of the platform with post-PNR simulations to analyze the relative power contributions of each module. This approach allowed us to estimate the static and dynamic power components for each mode at 0.8 V, as detailed in Section 3.2 and summarized in Table 1. The extraction of static power coefficients (P_S) and dynamic power coefficients (η) was crucial for our optimization, demonstrating the necessity of a meticulous characterization process.

4.2.2 FLL Timing Analysis and Modeling. In this section, we examine the timing overhead associated with FLL reconfiguration on the HEEPocrates platform, representing a typical FLL in RISC-V-based ULP platforms [4]. We then present a model for FLL configuration times, critical to our optimization strategy in Section 3 aimed at improving power efficiency for the case studies in Section 5.1.

In our analysis, we measured the FLL's transition times to lock onto various frequencies at the platform's minimum voltage of 0.8 V. Our findings indicate that transitions to frequencies closer to the current one incur less overhead than transitions to more distant frequencies. Moreover, transitions involving lower frequencies face increased overheads due to 1) longer register read and write times for actions like controller reprogramming, and 2) extended stabilization times at reduced FLL clock rates.

Given that we observed the shortest FLL reconfiguration times for frequencies slightly above 1 MHz, with times notably increasing

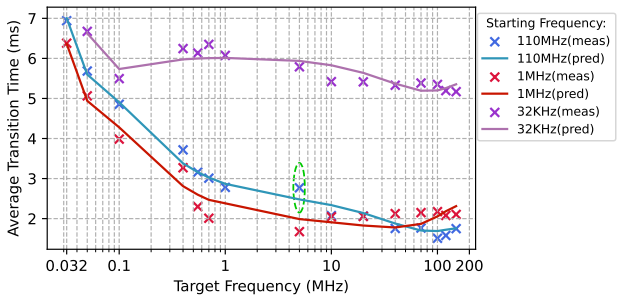

Figure 3: Model predictions vs. actual HEEPocrates FLL transition times. For example, the green point highlights the model-predicted and measured transition time for switching from 110 MHz to 5 MHz.

as frequencies decreased below this threshold, we designed our FLL model to employ both direct and reciprocal values of initial (F_1) and target (F_2) frequencies to represent the non-linear relationship across the frequency spectrum accurately. We used a polynomial model up to the third degree, incorporating squared, cubed, and interaction terms, effectively mapping frequency changes to configuration times. The accuracy of our third-degree polynomial model is evidenced in Figure 3, which compares the actual measured FLL transition times with the predictions of our model from various starting frequencies. The precision of the model, underscored by an R^2 score of 0.9817, attests to its efficacy in reliably forecasting FLL configuration times across various frequency transitions.

5 EXPERIMENTS

5.1 Optimal Frequencies for Our Case Studies

This section evaluates power efficiency after applying our proposed frequency optimization methodology to the biomedical applications introduced in Section 4.1, comparing the results against the conventional methodologies discussed in Section 2. Given the diversity of platforms used in existing studies, we adapt methods from the literature to HEEPocrates, considering an ADC buffer size of 256 samples. This comparison allows us to highlight the energy efficiency gains achieved through our optimization methodology.

In our comparative analysis, we revisit the primary approaches for managing the acquisition phase in wearables: the uniform All-On approach without specialized acquisition policies; the CS mode for energy-efficient continuous sampling; and the B_{lh} approach, which toggles between idle and active states at the platform's lowest and highest frequencies within the lowest voltage setting. For B_{lh}, we set frequencies to 32 kHz (idle) and 170 MHz (active) at 0.8 V, which corresponds to HEEPocrates' frequency range. For CS mode, to optimize dynamic power, we adjust operating frequencies to just exceed each application's bandwidth requirements. For example, the HBC_3 application's need for three ECG signals from 256 Hz results in setting the frequency to approximately 38 kHz after applying a compensation factor ($\lambda = 50$) for SPI communication efficiency, minimizing their dynamic power consumption.

We explore two configurations to optimize frequencies for burst mode, with the aim of minimal power consumption: (1) **uniform frequency**, setting $F_a = F_i$ and eliminating the FLL configuration time ($t_f = 0$), resulting in an optimal uniform frequency between idle and active states (referred to as Buni, detailed in Table 2), based on Equation 6; and (2) **dynamic frequency**, varying frequencies between idle (F_i) and active (F_a) states, incorporating the FLL configuration time model (Section 4.2.2) to optimize Equation 4 (identified as Bdyn, with optimized frequencies and their power consumption in Table 2).

Table 2 illustrates the power improvement achieved through our proposed frequency optimization strategies. In particular, the B_{dyn} configuration demonstrates superior power efficiency for the HBC_1, HBC_3, and CWM_4 applications, benefiting from the dynamic adjustment of frequencies between the idle and active states. This aligns with the FLL reconfiguration to optimize energy use. Conversely, for the HBC_12 application, which demands higher bandwidth, the B_{uni} configuration emerges as more efficient, highlighting the advantage of maintaining a uniform frequency for high-bandwidth applications. This observation underscores the adaptability of our approach to application-specific demands. The table also presents the frequencies corresponding to the minimized power consumption, allowing us to observe the direct impact of our optimization on the operational frequencies.

5.2 Total Energy Consumption Improvement

To understand the impact of our optimized frequency settings on energy consumption, considering both acquisition and processing times is crucial. The optimal frequency is determined by selecting the lowest voltage that enables a frequency where the system can meet its deadlines, then picking the highest possible frequency at that voltage to enable the fastest return to sleep mode (low-power acquisition phase). Figure 4 shows HBC_12's total energy consumption across processing frequencies, employing an optimal setting during acquisition, Buni. Thus, within their operational windows (15 s for HBC apps and 56 s for CWM_4), HBC_1, HBC_3, HBC_12, and CWM_4 execute in 22, 60, 210, and 2612 ms, respectively, at 170 MHz, each with a power of 4322.5 µW on HEEPocrates.

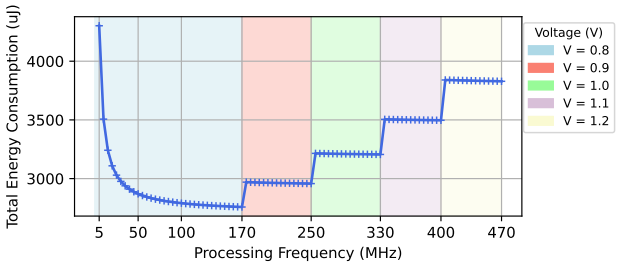

Figure 4: Total energy consumption for HBC_12 in one window, varying processing frequency, using Buni for acquisition.

We evaluated each application's energy consumption over a single operational window with varying acquisition strategies. Figure 5 presents the normalized energy consumption per second, dividing the energy of a single window by the duration of the window, allowing standardized comparisons of improvements in energy efficiency in applications with different window times.

Figure 5 delineates the energy distribution between the processing and acquisition phases of each methodology. It visually emphasizes the efficiency gains from our optimization, with the

Table 2: Comparison of acquisition power consumption and selected frequencies for HBC_1, HBC_3, HBC_12, and CWM_4 applications between our proposed method and methods found in the literature.

Method	Applications							
	HBC_1		HBC_3		HBC_12		CWM_4	
	Power (uW)	Freq (MHz)	Power (uW)	Freq (MHz)	Power (uW)	Freq (MHz)	Power (uW)	Freq (MHz)
B_{uni}	120.1	0.51	121.9	0.89	**126.4**	1.78	122.6	1.03
B_{dyn}	**118.8**	0.08, 3.53	**120.6**	0.10, 3.64	128.3	0.13, 3.74	**121.4**	0.11, 3.67
B_{lh}[1]	140.1	0.032, 170	184.5	0.032, 170	384.6	0.032, 170	206.8	0.032, 170
CS[2]	232.0	0.032[4]	232.1	0.038	233.3	0.15	232.2	0.05
All-On[3]	4322.5	170	4322.5	170	4322.5	170	4322.5	170

[1] B_{lh} approach as used in [1, 5]. [2] CS mode adopted in [14, 20, 21]. [3] All-On method utilized in [7, 8, 11–13].
[4] The lowest supported frequency in HEEPocrates.

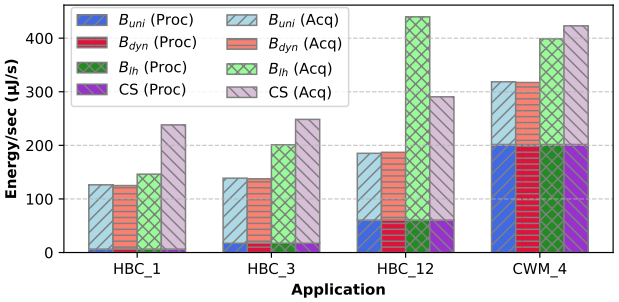

Figure 5: Normalized energy consumption for case study applications using different techniques, highlighting the proportion of energy allocated to acquisition and processing phases.

All-On approach excluded because of its disproportionately high energy consumption. In all of our case studies, our optimized frequencies lead to substantial energy savings. Specifically, for HBC_1, HBC_3, and CWM_4 applications, the dynamic frequency approach (B_{dyn}) consistently outperforms other methods, showcasing energy savings of 14.5 %, 31.7 %, and 20.4 % over the B_{lh} method, and 47.5 %, 44.7 %, and 25.0 % over the CS method, respectively. Meanwhile, for the high-bandwidth HBC_12 application, maintaining a uniform frequency (B_{uni}) is found to be the most effective, yielding a 57.9 % and 36.3 % reduction in total energy consumption compared to the B_{lh} and CS approaches.

6 CONCLUSIONS

In this study, we have introduced a novel approach for optimizing power during the biosignal acquisition phase in wearables by selecting optimal frequencies for low-power acquisition states. By modeling the system and analyzing key parameters, we have formulated the acquisition power based on target frequencies. Our analysis identifies optimal frequencies for both uniform and dynamic frequency selection between idle and active states. This includes variations where overheads depend on target frequencies, are fixed, or are negligible, catering to a wide range of operational contexts. We characterize our case study platform, detailing the FLL modeling and power coefficients for each state, which underpinned our optimization. Our methodology, applied to different biomedical applications, demonstrated notable energy savings and underscored the potential to extend the battery life of wearable devices.

ACKNOWLEDGMENT

This research is carried out in the frame of the "UrbanTwin: An urban digital twin for climate action: Assessing policies and solutions for energy, water and infrastructure" project with the financial support of the ETH-Domain Joint Initiative program in the Strategic Area Energy, Climate and Sustainable Environment. Also, this work was partly supported by a Ph.D. Student Grant for ESL by Nespresso S.A. within the EPFL-Nestlé Framework Agreement (EPFL TTO SoW no. 2019-0288).

REFERENCES

[1] R. Aghazadeh et al. 2018. Compressed sensing based seizure detection for an ultra low power multi-core architecture. In *HPCS*.
[2] Analog Devices. 2019. MAX86150. https://www.analog.com/en/products/max86150.html. Accessed: 7 March 2024.
[3] Analog Devices. 2023. AD4130-8. https://www.analog.com/en/products/ad4130-8.html. Accessed: 7 March 2024.
[4] D. Bellasi et al. 2017. A wide tuning-range ADFLL for mW-SoCs with dithering-enhanced accuracy in 65 nm CMOS. In *ISCAS*.
[5] S. Benatti et al. 2019. Online learning and classification of EMG-based gestures on a parallel ultra-low power platform using hyperdimensional computing. *TBioCAS* (2019).
[6] E. De Giovanni et al. 2020. Modular design and optimization of biomedical applications for ultralow power heterogeneous platforms. *TCAD* (2020).
[7] R. Donati et al. 2022. BioWolf16: A 16-channel, 24-bit, 4kSPS Ultra-Low Power Platform for Wearable Clinical-grade Bio-potential Parallel Processing and Streaming. In *EMBC*.
[8] S. Frey et al. 2023. BioGAP: A 10-core FP-capable ultra-low power IoT processor, with medical-grade AFE and BLE connectivity for wearable biosignal processing. In *COINS*.
[9] GreenWaves Technologies. 2021. GAP8. https://greenwaves-technologies.com/gap8_mcu_ai/. Accessed: 7 March 2024.
[10] GreenWaves Technologies. 2022. GAP9. https://greenwaves-technologies.com/gap9_processor/. Accessed: 7 March 2024.
[11] T. M. Ingolfsson et al. 2021. Towards long-term non-invasive monitoring for epilepsy via wearable eeg devices. In *BioCAS*.
[12] V. Kartsch et al. 2019. Biowolf: A sub-10-mw 8-channel advanced brain–computer interface platform with a nine-core processor and ble connectivity. *TBioCAS* (2019).
[13] V. Kartsch et al. 2019. An Energy-Efficient IoT node for HMI applications based on an ultra-low power Multicore Processor. In *SAS*.
[14] V. Kartsch et al. 2019. Ultra low-power drowsiness detection system with BioWolf. In *NER*.
[15] T. Lu et al. 2023. Biocompatible and long-term monitoring strategies of wearable, ingestible and implantable biosensors: reform the next generation healthcare. *Sensors* (2023).
[16] S. Machetti et al. 2024. X-HEEP: An Open-Source, Configurable and Extendible RISC-V Microcontroller for the Exploration of Ultra-Low-Power Edge Accelerators. *arXiv:2401.05548* (2024).
[17] L. L. Ng et al. 2022. Power Consumption in CMOS Circuits. In *Electromagnetic Field in Advancing Science and Technology*. IntechOpen.
[18] A. Pullini et al. 2019. Mr. Wolf: An energy-precision scalable parallel ultra low power SoC for IoT edge processing. *JSSC* (2019).
[19] R. Zanetti et al. 2021. Real-time EEG-based cognitive workload monitoring on wearable devices. *TBME* (2021).
[20] M. Zanghieri et al. 2019. Robust real-time embedded EMG recognition framework using temporal convolutional networks on a multicore IoT processor. *TBioCAS* (2019).
[21] S. Zanoli et al. 2023. An Error-Based Approximation Sensing Circuit for Event-Triggered Low-Power Wearable Sensors. *JETCAS* (2023).