

Alma Mater Studiorum – Università di Bologna

DOTTORATO DI RICERCA IN  
INGEGNERIA ELETTRONICA, TELECOMUNICAZIONI E  
TECNOLOGIE DELL'INFORMAZIONE

Ciclo XXXII

**Settore Concorsuale: 09/E3 - ELETTRONICA**

**Settore Scientifico Disciplinare: ING-INF/01 - ELETTRONICA**

Low-Power Human-Machine Interfaces: Analysis And Design

**Presentata da: VICTOR JAVIER KARTSCH MORINIGO**

**Coordinatore Dottorato**

**Supervisore**

**Prof. COSTANZO ALESSANDRA**

**Prof. BENINI LUCA**

**Esame finale anno 2020**

ALMA MATER STUDIORUM - UNIVERSITY OF BOLOGNA

# Low-Power Human-Machine Interfaces: Analysis And Design

by

Victor Javier Kartsch Morinigo

A thesis submitted for the degree of  
Doctor of Philosophy

in the  
Faculty of Engineering  
Department of Electrical, Electronic and Information Engineering "G. Marconi  
(DEI)

February 2020



*“Those who can imagine anything, can create the impossible.”*

Alan Turing



# *Abstract*

Faculty of Engineering

Department of Electrical, Electronic and Information Engineering "G. Marconi (DEI)

Doctor of Philosophy

by Victor Javier Kartsch Morinigo

Human-Machine Interaction (HMI) systems, once used for clinical applications, have recently reached a broader set of scenarios, such as industrial, gaming, learning, and health tracking thanks to advancements in Digital Signal Processing (DSP) and Machine Learning (ML) techniques. A growing trend is to integrate computational capabilities into wearable devices to reduce power consumption associated with wireless data transfer while providing a natural and unobtrusive way of interaction. However, current platforms can barely cope with the computational complexity introduced by the required feature extraction and classification algorithms without compromising the battery life and the overall intrusiveness of the system. Thus, highly-wearable and real-time HMIs are yet to be introduced.

Designing and implementing highly energy-efficient biosignal devices demands a fine-tuning to meet the constraints typically required in everyday scenarios. This thesis work tackles these challenges in specific case studies, devising solutions based on bioelectrical signals, namely EEG and EMG, for advanced hand gesture recognition.

The implementation of these systems followed a complete analysis to reduce the overall intrusiveness of the system through sensor design and miniaturization of the hardware implementation. Several solutions have been studied to cope with the computational complexity of the DSP algorithms, including commercial single-core and open-source Parallel Ultra Low Power architectures, that have been selected accordingly also to reduce the overall system power consumption. By further adding energy harvesting techniques combined with the firmware and hardware optimization, the systems achieved self-sustainable operation or a significant boost in battery life.

The HMI platforms presented are entirely programmable and provide computational power to satisfy the requirements of the studies applications while employing only a fraction of the CPU resources, giving the perspective of further application more advanced paradigms for the next generation of real-time embedded biosignal processing.



## *Acknowledgements*

We are the product of the effort of uncountable persons that have put on us hope on creating something beautiful. To this, in the first term, I would like to thank my family. I want to thank also to Prof. Luca Benini and Dr. Simone Benatti that have welcomed me and gave me the chance to grow and learn, to Dr. Marco Guermandi for his unconditional help and patience, and to Prof. Davide Rossi for his invaluable support. This experience could not be as remarkable as it is without my friends and colleges from the lab, Fabio, Mattia, Alberto, and Tommaso, thank you all. Finally, I thank you, Mayra, to be my everyday companion, a light that guided me in harsh moments and completed me in moments of happiness.





# Contents

<b>Abstract</b>	<b>iv</b>
<b>Acknowledgements</b>	<b>vi</b>
<b>List of Figures</b>	<b>xi</b>
<b>List of Tables</b>	<b>xiii</b>
<b>Abbreviations</b>	<b>xiv</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Thesis contribution . . . . .	2
1.2 Thesis structure . . . . .	3
<b>2 Acquisition and Processing of Biosignals for HMI</b>	<b>4</b>
2.1 Biosignals for Human-Machine Interaction . . . . .	4
2.1.1 EMG Signals . . . . .	5
2.1.2 EEG Signals . . . . .	5
2.2 Biopotential Signal Acquisition . . . . .	8
2.3 Processing Unit . . . . .	10
2.3.1 ARM Cortex-M4 SoC . . . . .	10
2.3.2 PULP platform . . . . .	11
2.3.2.1 PulpV3 SoC . . . . .	11
2.3.2.2 Mr. Wolf SoC . . . . .	12
2.4 Signal Processing . . . . .	12
2.4.1 Preprocessing . . . . .	13
2.4.2 Feature Extraction . . . . .	14
2.4.3 Classification . . . . .	15
<b>3 Hardware Design and Implementation of HMIs</b>	<b>17</b>
3.1 ARM-based HMI . . . . .	17
3.1.1 PCB design . . . . .	18
3.1.2 Sensor Interface . . . . .	18
3.1.3 Subsystems . . . . .	19
3.2 BioWolf: The Parallel Ultra Low Power HMI . . . . .	19

3.2.1	PCB design . . . . .	20
3.2.2	Sensor Interface and electrical characterization . . . . .	21
3.2.3	Operational Modes . . . . .	23
3.2.4	Subsystems . . . . .	24
3.2.5	Energy Harvesting . . . . .	24
3.2.6	Software . . . . .	25
3.2.7	Discussion . . . . .	26
<b>4</b>	<b>HMI applications using BioWolf</b>	<b>28</b>
4.1	Drowsiness Detection . . . . .	28
4.1.1	Introduction . . . . .	29
4.1.2	Feature Extraction . . . . .	30
4.1.3	Sensor Fusion . . . . .	31
4.1.4	Embedded Deployment . . . . .	33
4.1.5	Experimental Results . . . . .	34
4.1.6	Embedded System Performance . . . . .	35
4.1.7	Implementation on the PULP architecture . . . . .	36
4.1.8	Drowsiness Detection With BioWolf . . . . .	39
4.1.8.1	Feature Extraction . . . . .	39
4.1.8.2	Training and Classification . . . . .	40
4.1.8.3	Experimental Results . . . . .	41
4.1.9	Discussion . . . . .	42
4.2	SSVEP-Based BCI . . . . .	42
4.2.1	Introduction . . . . .	43
4.2.2	BCI System . . . . .	44
4.2.3	Signal Processing . . . . .	46
4.2.4	Parallel Processing and Optimizations . . . . .	48
4.2.5	Experimental Results . . . . .	50
4.2.6	Power Consumption . . . . .	51
4.2.7	Discussion . . . . .	52
4.3	ERP-Based BCI . . . . .	54
4.3.1	Introduction . . . . .	54
4.3.2	Signal Processing . . . . .	55
4.3.3	Experimental Results . . . . .	56
4.3.4	Discussion . . . . .	59
4.4	Online Learning and Classification of EMG-Based Gestures using Hyper-dimensional Computing . . . . .	59
4.4.1	Introduction . . . . .	59
4.4.2	Signal Processing and Classification . . . . .	61
4.4.3	Experimental Results . . . . .	62
4.4.4	Discussion . . . . .	65
<b>5</b>	<b>Conclusions</b>	<b>68</b>

**Bibliography**

**70**

# List of Figures

2.1	Typical components of an HMI system, including the signal source, the acquisition device, and the application. . . . .	5
2.2	Raw EMG signal during several contractions and rest muscular cycles. . .	6
2.3	EEG response (time-domain) with eyes closed. . . . .	6
2.4	Brain rhythms obtained by filtering the raw EEG signal into the different corresponding bandwidths. . . . .	7
2.5	EEG P300 response obtained after averaging several EEG trials. . . . .	8
2.6	Electrode-skin equivalent circuit. From the electrical standpoint, the skin behaves as several resistors and capacitors connected in series (and parallel) to the acquisition device. . . . .	9
2.7	Mr. Wolf SoC architecture including the SoC and Cluster domains. . . . .	13
2.8	Typical Signal Processing Chain for HMIs. In an embedded system, all computation is performed in the sensor node. . . . .	13
2.9	Example of the signal output after feature extraction of EMG signals. The top-left figure presents the original signal while the remaining present the different signal extraction outputs for RMS, Waveform Length, and Discrete Wavelet Transform. . . . .	15
3.1	ARM BCI including Block diagram and PCB implementation. . . . .	18
3.2	Scheme of the electrode configuration used for both systems*. . . . .	19
3.3	BioWolf System. Block diagram (left) and PCB implementation (left). . .	20
3.4	Example of BioWolf montage in an SSVEP-based BCI for hands-free control of a digital score for musicians. . . . .	21
3.5	Operational modes of BioWolf. . . . .	24
3.6	Solar Panel current charging output for different illumination conditions. .	25
3.7	BioWolf GUI for streaming and data storing (PC). . . . .	26
3.8	BioWolf GUI for Android during EMG classification. . . . .	26
4.1	Features extracted from EEG and IMU sensors. . . . .	30
4.2	Block diagram of the extracted features to compose the drowsiness sensor fusion algorithm. . . . .	33
4.3	Block diagram of the complete wearable deployment for drowsiness detection based on the ARM Cortex-M4 MCU. . . . .	33
4.4	Detection accuracy of the different drowsiness alarms. . . . .	34
4.5	Power consumption of the embedded implementation (on PCB). The MCU dominates the power consumption, employing 87% of the total power. .	35
4.6	Current drawn by the MCU during the processing and classification of the drowsiness levels. . . . .	36
4.7	Power consumption comparison between different platforms. . . . .	38

4.8	Block diagram of the feature extraction and classification for the drowsiness detection with BioWolf. . . . .	39
4.9	Average CCA correlation of SSVEP responses for different stimuli (x-axes) calculated with different reference signals (y-axes). . . . .	44
4.10	BCI setup during experimentation. . . . .	45
4.11	Block diagram of CCA algorithm and implementation. . . . .	46
4.12	Average ITR results calculated with different thresholds values. . . . .	47
4.13	Diagram of the signal sampling and execution of the application on Mr. Wolf. . . . .	47
4.14	Parallelization schema of the processing pipeline highlighting the execution flow and its branches. This diagram reports the size of data samples and the parallelization approach (loop-level or task-level) for each computing block. . . . .	49
4.15	BCI testing after one hour of usage. A - CCA test, where the solid lines denote the CCA output of the SSVEP response on a test performed minutes after setting up the BCI system. The dashed lines denote the CCA output of the SSVEP response after one hour of device wear. B - Alpha waves test. The topmost figure of this section presents the alpha band power of a 55-second trial, where the blue line shows the PSD response form EEG signals acquired on a control test minutes after setting up the device, while the red line corresponds to a test after one hour of device wear. The bottom figure denotes the signal in the time-domain of the highlighted (in red) region of the first figure. . . . .	53
4.16	Setup and processing steps for the ERP-based BCI. . . . .	57
4.17	ERP from subject 1. . . . .	57
4.18	Grand average of the ERPs extracted from the five subjects. . . . .	58
4.19	Implementation on BioWolf of the HD computing algorithm. . . . .	62
4.20	Gestures used for the testing. Open hand, fist, index, 2-fingers pinch, rest position. . . . .	63
4.21	Average accuracy obtaining by HD computing, using the same data collected by 10 subjects, increasing the number of gestures (from 1 to 11). . . . .	63
4.22	Battery duration of the target applications including the static power of the ADC with an increasing number of channels. . . . .	66

# List of Tables

2.1	Mr.Wolf SoC features. . . . .	12
2.2	Typical Feature Extraction Techniques for Biosignals . . . . .	14
3.1	Relevant features of the devices used in the system. . . . .	20
3.2	Contact impedance measurement for different values of $Z_{C1,2}$ and $Z_{IN}$ . . . . .	23
3.3	Comparison between SoA embedded platforms for BCI processing. . . . .	27
4.1	Normalized peak-to-peak difference for different blinks and detection accuracy. . . . .	32
4.2	Threshold values for ST, NOM and NM calculated from the IMU sensor. . . . .	32
4.3	Number of cycles required to compute each function in the analyzed embedded computing platforms [KCycles]. . . . .	36
4.4	Comparison between different platforms for real-time operation. . . . .	38
4.5	System Drowsiness Levels . . . . .	40
4.6	K-means Offline test per level . . . . .	41
4.7	Online accuracy per level for the NCC . . . . .	41
4.8	Online results for 4 stimuli. . . . .	50
4.9	Energy/Classification per Kernel Function . . . . .	51
4.11	SoA embedded implementations for SSVEP classification. . . . .	54
4.12	HD Computing Execution times on the target architectures, with 10,000-D, N=1. (Cyc, su) stand for (cycles, speed-up). The total energy/class reported, is the result of the addition of the contribution of these functions without considering the energy during idle periods. . . . .	64
4.13	Current Consumption of the Board Components in the Different Operational States . . . . .	65
4.14	Comparison with SoA systems for EMG classification. . . . .	66

# Abbreviations

<b>ADC</b>	<b>A</b> nalog-to- <b>D</b> igital <b>C</b> onverter
<b>AFE</b>	<b>A</b> nalog <b>F</b> ront <b>E</b> nd
<b>ANN</b>	<b>A</b> rtificial <b>N</b> eural <b>N</b> etwork
<b>BCI</b>	<b>B</b> rain <b>C</b> omputer <b>I</b> nterface
<b>BMI</b>	<b>B</b> rain <b>M</b> achine <b>I</b> nterface
<b>CCA</b>	<b>C</b> anonical <b>C</b> orrelation <b>A</b> nalysis
<b>DSP</b>	<b>D</b> igital <b>S</b> ignal <b>P</b> rocessing
<b>ECoG</b>	<b>E</b> lectro <b>C</b> ortico <b>G</b> raphy
<b>EEG</b>	<b>E</b> lectro <b>E</b> ncephalo <b>G</b> raphy
<b>EMG</b>	<b>E</b> lectro <b>M</b> yo <b>G</b> ram
<b>EOG</b>	<b>E</b> lectro <b>O</b> culo <b>G</b> raphy
<b>FFT</b>	<b>F</b> ast <b>F</b> ourier <b>T</b> ransform
<b>FPU</b>	<b>F</b> loating <b>P</b> oint <b>U</b> nit
<b>FTA</b>	<b>F</b> requency <b>T</b> agging <b>A</b> nalysis
<b>HCI</b>	<b>H</b> uman <b>C</b> omputer <b>I</b> nterface
<b>HMI</b>	<b>H</b> uman <b>M</b> achine <b>I</b> nterface
<b>GUI</b>	<b>G</b> raphic <b>U</b> ser <b>I</b> nterface
<b>IIEG</b>	<b>I</b> ntracranial <b>E</b> EG
<b>IoT</b>	<b>I</b> nternet <b>o</b> f <b>T</b> hings
<b>ITR</b>	<b>I</b> nformation <b>T</b> ransfer <b>R</b> ate
<b>PCA</b>	<b>P</b> rincipal <b>C</b> omponent <b>A</b> nalysis
<b>PCB</b>	<b>P</b> rinted <b>C</b> ircuit <b>B</b> oard
<b>PLI</b>	<b>P</b> ower <b>L</b> ine <b>I</b> nterference
<b>PSD</b>	<b>P</b> ower <b>S</b> pectral <b>D</b> ensity <b>A</b> nalysis
<b>SPI</b>	<b>S</b> erial <b>P</b> eripheral <b>I</b> nterface



<b>SVD</b>	<b>S</b> ingular <b>V</b> alue <b>D</b> ecomposition
<b>SVM</b>	<b>S</b> upport <b>V</b> ector <b>M</b> achine
<b>SSVEP</b>	<b>S</b> tady <b>S</b> tate <b>V</b> EP
<b>SNR</b>	<b>S</b> ignal-to- <b>N</b> oise <b>R</b> atio
<b>VEP</b>	<b>V</b> isual <b>E</b> voked <b>P</b> otential

# Chapter 1

## Introduction

Embedded wearable applications have gained a considerable boost in popularity, driven by the development of unintrusive and low-power digital systems. As this happens, biopotential monitoring devices are gaining traction considering their potential not only in health applications but also to anticipate or infer human intentions in everyday's life on healthy subjects.

Among the plethora of information available from the human body, neural activity (EEG) has attracted significant interest since it holds the promise to provide a direct brain-computer link [1] [2] with limited invasivity. As a consequence, several commercial platforms, such as Emotiv Insight and EPOC+ [3], Neurosky MindWave [4], and OpenBCI [5] already provide simple BCIs features, enabling new application scenarios beyond the biomedical, such as industrial and gaming[6].

Furthermore, hand gesture recognition is a critical element to enable natural and advanced ways of communication between objects and users in many domains, in the track of the ever-evolving IoT trend. One of the major approaches used for hand gesture recognition, also in commercial scenarios, is the processing of muscular activity [7] through electromyography (EMG). Prosthetic systems such as [8] and [9] can decode predefined bursts of muscular contractions to produce a highly reliable and amenable wearable system. An approach that is gaining popularity is based on machine learning (ML) techniques to analyze EMG signal patterns during muscular contractions.

Embedding these systems is crucial to enable compelling human-machine interaction. The options explored to embedding BCIs processing include using portable single-board computers, streaming acquired data to an external device, and embedding acquisition and processing on a single device. The latter, also known as edge computing,

is considered a more energy-efficient approach as it reduces the power intrinsic to data transmissions.

Nevertheless, porting applications, for both EEG and EMG, to wearable and embedded systems is not a trivial task, posing several challenges in terms of computational power, battery life, and wearability. For instance, feature extraction and classification has to be supported by extremely high-efficiency hardware and firmware optimization to satisfy real-time constraints and minimize power consumption simultaneously.

Clinical-grade EEG/EMG systems are based on cumbersome systems that are hardly socially acceptable in a non-clinical environment [10]. And, although some of the commercial platforms listed above provide a less intrusive interface, they lack on signal quality [11]. Likewise, this inconvenient is found for the processing and acquisition node. Commercial and research nodes with a practical form factor have yet to be delivered [5, 12].

The development of a complete, low-intrusive, fully wearable device is still an open challenge. Designing and implementing highly energy-efficient biosignal devices requires a multi-level approach to balance the challenges imposed by everyday-like devices. This work offers a complete analysis to reduce the overall intrusiveness of the system through sensor design and miniaturization of the hardware implementation and the use of both commercial (ARM) and open-source (PULP) MCU architectures as to provide energy-efficient and real-time processing. The next section describes the specific contribution of this thesis.

## 1.1 Thesis contribution

This work tackles the previously introduced challenges from many perspectives, including both hardware and firmware optimizations, in an application-oriented fashion. Significant effort has been granted to the develop of BioWolf, a multicore highly configurable platform that meets most of the modern HMIs requirements for everyday scenarios. The development of BioWolf aims to provide solutions in terms of intrusiveness, performance, and energy efficiency. Intrusiveness has been significantly reduced through improvements in the sensor interface and PCB miniaturization, obtaining an area reduction of 4x to state-of-the-art (SoA) systems. In terms of energy efficiency, BioWolf can provide 4x more battery life than single-core commercial base platforms while only employing a fraction of the available processing power (<20%) and achieving State-of-the-art (SoA) accuracy for both EEG and EMG applications (>90% for most cases). Furthermore, by adding energy harvesting techniques combined with the

firmware and hardware optimization, some of the presented systems can achieve self-sustainable operation or a significant boost in battery life (2-4x).

In the following sections, details of BioWolf design and its features are introduced. Consequently, four study cases of biosignal classification and embedded implementation are presented. Concretely, this work focuses on three study cases on EEG classification, including a drowsiness detection system, an SSVEP-based BCI, an ERP-based BCI and, finally, an EMG system for Gesture Recognition.

The results demonstrate the benefits of multicore microcontroller architectures for biomedical signals for HMI to commercially available systems based on single-core processors, empowering the next generation of real-time embedded biosignal processing.

## **1.2 Thesis structure**

The rest of the thesis is structured as follows.

The second chapter, named, Acquisition and Processing of Biosignals for HMI, provides background information about the HMIs, describing in detail all the modules that compose a typical HMI system and the challenges associated with a wearable deployment. Similarly, this chapter introduces the signal processing and classification techniques to overcome the problems related to the stochastic and non-stationary behavior of biosignals while being suitable for embedded implementation.

The third chapter, namely, Hardware Design, and Implementation of HMIs, provides a summary and experimental results on the HMI architectures introduced in this work, giving particular emphasis to BioWolf as it is the main focus of this thesis.

The fourth chapter, HMI applications using BioWolf, introduces details about the signal processing, classification, embedded deployment, and optimizations of four case studies for HMI. The first three (EEG) include a drowsiness detection system, a CCA-based BCI for SSVEPs, and a system for food quality grading through ERPs, while the fourth proposes an EMG gesture recognition system based on High-Dimensional Computing. Each section offers a discussion to highlight the improvements achieved for each application by employing the resources available in BioWolf to previous work and the SoA.

Finally, chapter five draws the final conclusions of this thesis work.

## Chapter 2

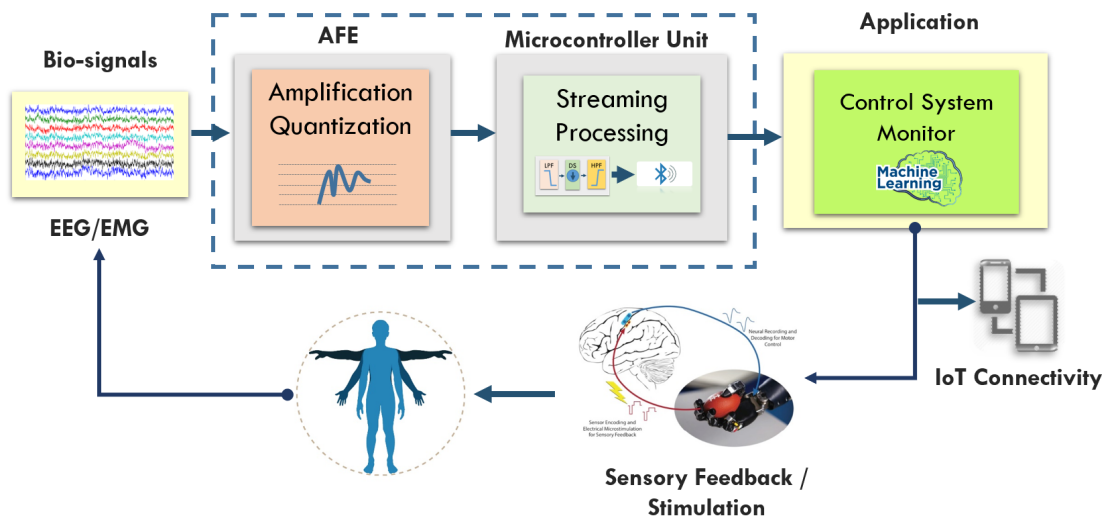
# Acquisition and Processing of Biosignals for HMI

Biosignals contain qualitative and quantitative information about the human body's state, which has been exploited in the past to improve the diagnosis and treatment of health disorders. Most recently, it has also gained growing popularity to enhance interaction with the external world through Human-Machine Interaction devices.

Fig. 2.1 highlights the main building blocks of an HMI system, which includes the biosignal source, the analog-front-end, and the processing unit. A specific application then tailors the resulting information to signal processing and classification techniques to provide the control, monitoring, or feedback. This chapter offers a detailed introduction about these blocks highlighting the challenges to provide unobtrusive and energy-efficient HMIs.

### 2.1 Biosignals for Human-Machine Interaction

The body can provide a variety of signal sources that include, for instance, electrocardiogram (ECG), electrooculogram (EOG), and galvanic skin response (GSR). Nonetheless, those above are not suitable for HMI as not under volitional control from the user, and therefore, only useful for monitoring. On the other hand, Electroencephalography (EEG), EMG, and ECoG can be used for biomedical control and HMI, therefore, providing a broader range of applications. Nevertheless, only EMG and EEG are practical for everyday scenarios. For these reasons, EEG and EMG are the main focus of this work, and the following aims to provide insights about the body generation and modulation mechanisms.



**Figure 2.1:** Typical components of an HMI system, including the signal source, the acquisition device, and the application.


### 2.1.1 EMG Signals

EMG signals are non-stationary stochastic electric potentials originated by the current flow of ionic charges through the membrane of the muscular fibers. Muscle activation is triggered by the central nervous system through signals sent to the motor units (motor neurons), innervating the muscular tissue. Each motor neuron is associated with neuromuscular junctions and muscle fibers through its axon. The activation of the muscle fibers depends on the motor neurons depolarization threshold and its size [13]. The final EMG signal is thus, the sum of the different motor units.


EMG can be collected using invasive or non-invasive methods[14]. Invasive methods consist of inserting recording needles to a specific muscle, allowing an excellent characterization of the EMG response even for deeper fibers. Non-invasive techniques instead capture the muscle (sEMG) activation, only requiring electrodes attached to the skin. However, this method does not allow for fine characterization of the EMG source since it only captures the surface activations. Still, sEMGs can be used in a variety of applications, including medical (muscular disease) and HMI. Fig. 2.2 presents a typical sEMG response for muscle contraction. When performing different hand gestures, the signal increase in amplitude with respect to the rest states.

### 2.1.2 EEG Signals

EEG signals are electric potentials generated on the scalp during the ionic transfer of large groups of pyramidal neurons in the cerebral cortex. The neurons receive and transmit electrochemical signals through axons and dendrites. The axons transport



**Figure 2.2:** Raw EMG signal during several contractions and rest muscular cycles.




**Figure 2.3:** EEG response (time-domain) with eyes closed.

electrical impulses while the dendrites bring information to the cell bodies. The ionic charges are due to interactions between  $\text{Na}^+$ ,  $\text{Ca}^{++}$ ,  $\text{K}^+$  cations, and  $\text{Cl}^-$  anions. The registered electric potential is the result of the summation of the individual potentials of large groups of neurons firing synchronously [15], which results in a stochastic and non-stationary process [16]. Fig. 2.3 shows a typical EEG response to the closure of the eyes.

Brain signals have been of great interest from early times since they hold the promise of providing not only brain function information but also the status of the whole body [17]. It has been used for diagnosis and treatment of brain disorders, while also attracting significant interest in non-medical scenarios such as in industry and gaming [6].

EEG analysis is mainly based on the spectrum of the naturally occurring activity or during auditory or visual stimulation. The spectral components of the EEG can be separated into bands (waves or rhythms), called *delta*, *theta*, *alpha*, *beta*, and *gamma*. Delta waves (0.5-4Hz) are visible during deep sleep stages (3rd-NREM) and walking.




**Figure 2.4:** Brain rhythms obtained by filtering the raw EEG signal into the different corresponding bandwidths.

[18]. Theta waves activity (4-7.5Hz) have been linked to several causes, including deep meditation, level of arousal, sleep stage transitions [19, 20]. Alpha activity (8.5-13Hz) is associated with relaxed awareness, and it is visible, with high amplitude (highest among all the rhythms), with the eyes closed [21]. Beta activity (14-26Hz) is often associated with active thinking (focus, attention, and problem solving), anxious, or panic events. In the motor cortex, beta waves are also associated with muscle contractions [22, 23]. Gamma rhythms (>30Hz) are rarely present on the EEG, although its presence can be evidence of a brain disease [17]. Fig. 2.4 presents an example of the brain rhythms during wakeful resting. These rhythms are naturally present on the brain [24, 25], and they can be used to monitor the physiological parameters of the user, such as drowsiness.

Event-related potentials (ERP), on the other hand, are direct time-locked EEG response to a motor or sensory stimulus or a cognitive task. ERPs are obtained through signal averaging to increase the signal-to-noise ratio (SNR) as they are small in amplitude. The signals have positive and negative components (P and N) and are typically visible after a few milliseconds (>100ms) from the stimuli presentation (non-obligatory). Fig. 2.5 shows an example of ERPs.

The importance of ERPs is two-folded. First, it can be used for the diagnosis of brain disorders. For instance, P300 can be used to detect early stages of Alzheimer's disease [26] as its response is smaller to control groups. Secondly, it can be used as a way of HMI. In such a scenario, P300 is elicited by presenting a target stimulus (the flashing of the intended symbol) occasionally, among non-relevant ones (the other





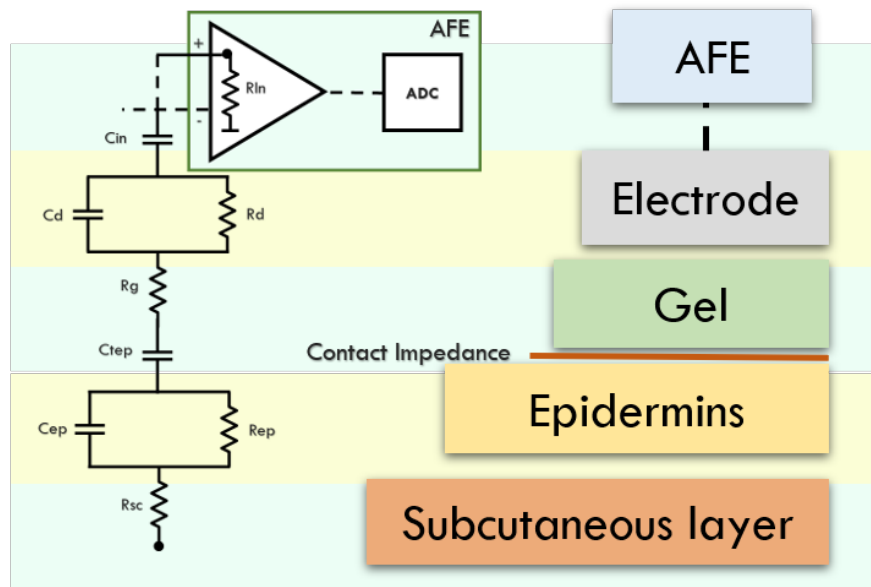
**Figure 2.5:** EEG P300 response obtained after averaging several EEG trials.

possible symbols), and can be used to detect the choice of a subject. Steady-State Visual Evoked Potentials (SSVEP) has also been used recently with considerable success [27–30] for HMI. SSVEP is elicited in the primary visual cortex as a result of a periodic visual stimulation and is, therefore, phase and frequency locked to the source. The detection is typically performed through spectral analysis, although recent work demonstrated that Canonical Correlation Analysis (CCA) [31] could lead to better performance for HMI.

## 2.2 Biopotential Signal Acquisition

Biosignals are particularly challenging due to their small amplitude, high impedance, and noise. Following strict requirements, signals need to be amplified to match the characteristics of the Analog-to-digital converter (ADC). Hence, the Analog-front-end (AFE) plays a crucial role to obtain information from a specific event, especially when the amplitude of the signal, like in the case of EEG, is small ( $<50\mu\text{V}$ ). Similarly, for clinical applications, the AFE must follow the International Federation of Clinical Neurophysiology (IFCN) standards [32], as they ensure good quality of the digital recording.

Fig. 2.6 presents a typical electrode-skin and AFE equivalent circuit that highlights the structure and challenges for the acquisition of biosignals. An essential factor that determines the quality of the signal is the contact impedance between the electrodes and the skin ( $Cd||Rd$ ). A high value of contact impedance leads to a potential divider effect at the input of the amplifier, which causes a reduction in Common-Mode Rejection Mode (CMRR) [33], hence allowing for common-mode noise, such as Power Line Interference (PLI) to be introduced in the system. Traditional systems employ saline solutions or



**Figure 2.6:** Electrode-skin equivalent circuit. From the electrical standpoint, the skin behaves as several resistors and capacitors connected in series (and parallel) to the acquisition device.

electroconductive gels to minimize these problems. Nevertheless, they are only suitable for laboratory conditions.

Zero-preparation dry electrodes, on the other hand, are suitable for use outside clinical settings due to the quick setup features, even for untrained users. However, they add some challenges such as a steep increase in contact impedance from the typical values of few tens of  $K\Omega$  for wet electrodes after extensive skin preparation, to hundreds of  $K\Omega$  or even several  $M\Omega$ . It also increases the noise generated at the metal-skin interface and augments the effect of interference coupling through capacitive effects on the cables or artifacts due to cable movement, microphony, and piezoelectric effects. To tackle these issues, the so-called active-electrode approach [34, 35] is typically used, and it has also been adopted in the present work.

As introduced above, once the signals are generated, signal conditioning is needed before quantization. An approach that has gained popularity is to use integrated circuits that provide both signal conditioning, sampling, and quantization. From the many solutions including ASIC [36, 37] and commercial [38, 39], the ADS1298, from Texas Instruments (TI), stands out. It is currently considered a standard for biopotential acquisition platforms due to its favorable trade-off between performance and power consumption ( $0.85\text{mW}/\text{Channel}$ ) while being compatible with the IFCN standards [32]. The AFE allows simultaneous sampling of up to eight bipolar channels, at up to 32 kbps sampling rate and 24-bit resolution. Each channel can be amplified by a factor 1 to 12 through an internal Programmable Gain Amplifier (PGA).

## 2.3 Processing Unit

Miniaturization of electronic devices such as the AFE allows for a paradigm shift from external data analysis and classification to wearable onboard processing. Hence, the computational unit plays a growing role in modern HMIs. Nevertheless, porting applications, for both EEG and EMG, to wearable and embedded systems is not a trivial task as the final system requires to provide both computational power (real-time constrains) and energy efficiency (battery life).

With this goal, we have focused first on commercially available platforms such as the ARM Cortex-M4, which provides a compatible power budget while delivering the required computational. A step forward has been achieved later through Mr. Wolf, a Parallel Ultra-low Power (PULP) architecture. This section introduces the general characteristics of both architectures.

### 2.3.1 ARM Cortex-M4 SoC

The ARM Cortex-M is a 32-bit RISC microcontroller that includes DSP instructions and a floating-point unit, making it suitable for biosignal processing. Several BCI systems have adopted this architecture [40–43], to allow for online processing.

Throughout this work, the STMicroelectronics STM32F407 has been used extensively. The processor is based on Harvard architecture with three-pipeline stage and branch speculation, implementing 32-bit RISC ISA. It also includes the entire 16-bit *Thumb*<sup>®</sup> – 1 and 16/32-bit *Thumb*<sup>®</sup> – 2 instruction sets. It features an FPU for fast floating-point operations and it is equipped with 192 KB of SRAM and 1 MB of non-volatile Flash memory. It features a power density of 238  $\mu\text{W}/\text{MHz}$ , offering up to 210 DMIPS and 566 CoreMark at 168 MHz on general-purpose applications.

Its rich set of peripheral (SPI, I2C, UART) allows flexibility to communicate with external devices such as specialized ADC for biopotentials. It features non-maskable interrupts(NMI) and up to 240 physical interrupts, with different priority levels (from 8 to 256). It also offers several power modes, allowing to reduce the overall power consumption of the system.

It also features a JTAG and Serial-Wire Debug (SWD) ports with up to eight breakpoints and four watchpoints. The processor also features Memory Protection Unit (MPU) for the eight memory sub-regions, preventing access to privileged application data.

### 2.3.2 PULP platform

PULP platform is an open-source project created by the collaboration of the University of Bologna and the Eth Zurich. PULP is a many-core platform (OpenRISC or RISC-V ISA) able to operate in broad range voltages and frequencies, which provides energy efficiency for both low and computationally heavy applications.

Several versions of the PULP architecture have been developed, and they differ on the chip technology (CMOS, FD-SOI implemented on 65, 40, or 28nm), the number of processing units, and memory size. In the following, PulpV3 and Mr. Wolf are introduced as they have been included in the systems and applications presented in this work.

#### 2.3.2.1 PulpV3 SoC

PulpV3 is a 28-nm UTBB FD-SOI SoC [44] that includes a cluster of four cores, where each processor implements a power-optimized four-pipeline stage micro-architecture based on the OpenRISC ISA [45].

The original OpenRISC ISA and the micro-architecture have been enhanced for energy-efficient DSP, supporting load and store operations embedding pointer arithmetic and power management instructions, and zero-overhead hardware loops.

The cluster of cores features a shared instruction cache with L0 buffer and support instruction broadcasting as to reduce the load on the cache banks[46]. The cluster relies on an L1 explicitly-managed multi-banked Tightly-Coupled Data Memory (TCDM) for energy efficiency and fast data access. A Direct Memory Access (DMA) unit is used to explicitly transfer data from the off-cluster 256kB L2 data memory and the TCDM (L1).

Off-cluster, several peripherals such as SPI and I2C communicate with the external world through a micro-DMA subsystem able to autonomously transfer data from the peripherals to the L2 memory while the cluster is idling, improving the energy efficiency of the system. The SoC and the cluster are in two different voltage and frequency domains to further enhance energy efficiency depending on the computational workload of applications. Moreover, an automatic clock gating mechanism is used to reduce the power consumption of idle resources of the system.

The PULP platform relies on OpenMP 3.0 parallel library that operates on top of a GCC 4.9 toolchain for programming. The hardware/software environment of PULP

includes a set of software tools useful to implement and debug applications that run on the architecture, and to estimate their execution time.

### 2.3.2.2 Mr. Wolf SoC

Mr. Wolf is a multi-core programmable SoC implemented in TSMC 40 nm CMOS technology, combining a 12 Kgates RISC-V processor for power and peripheral managing (Fabric Controller, FC) with a parallel cluster of eight RISC-V processors equipped with flexible and powerful DSP extensions, including floating-point units [47]. These blocks reside in two different power domains. Hence, the voltage and the frequency of these two components can be adjusted to meet the application's requirements with high energy efficiency.

The SoC features a full set of peripherals, including a JTAG interface for debug purposes, SPI, UART, and GPIOs, enabling parallel capture and transmission of bio-potentials at high bandwidth and efficiency. The peripherals, connected to the 512 kB L2 memory through an advanced I/O subsystem, reside on a third clock domain allowing adjusting the frequency of each peripheral according to the I/O requirements of applications. The peripheral subsystem has a dedicated DMA channel ( $\mu$ DMA) that autonomously controls the data transfer to/from the L2 memory. The  $\mu$ DMA has 2 dedicated 32-bit ports on the L2 memory subsystem, granting an aggregated bandwidth equal to  $2 \times 32$ -bit multiplied by the SoC clock frequency. Table 2.1 summarizes Mr. Wolf features while Fig. 2.7 presents a block diagram of the complete SoC.

## 2.4 Signal Processing

As introduced previously, the biosignals can be both processed online and offline. A typical signal processing scheme (in an embedded system) is presented in Fig.2.8, which

**Table 2.1:** Mr.Wolf SoC features.

Technology	CMOS 40nm LP
Chip Area	10mm <sup>2</sup>
Memory Transistors	576 kB
Equivalent Gates (NAND2)	1.8 Mgates
Voltage Range	0.8 V – 1.1 V
Frequency Range	32 kHz – 450 MHz
Sleep Power (State Retentive)	72 $\mu$ W (108 $\mu$ W)
SoC Power Density @ 0.8V	33 $\mu$ W/MHz
Cluster Power Density @ 0.8V	171 $\mu$ W/MHz

At its most efficient operating point (100 MHz, 0.8 V), Mr. Wolf can perform up to 15 Mega multiply-accumulate (MAC) operations per s/mW and 9 Mega Fused-MAC/s/mW.

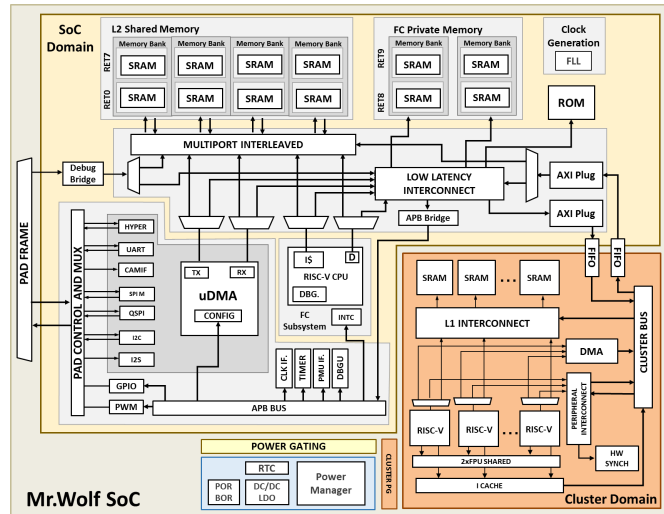


Figure 2.7: Mr. Wolf SoC architecture including the SoC and Cluster domains.

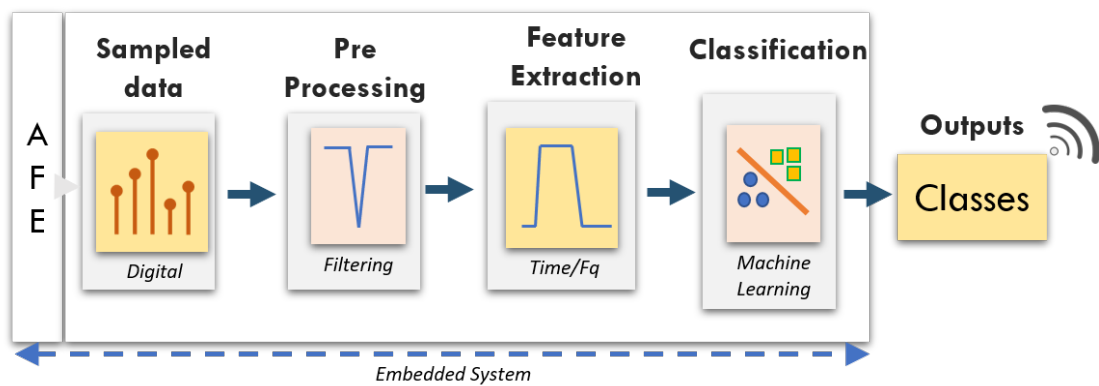


Figure 2.8: Typical Signal Processing Chain for HMIs. In an embedded system, all computation is performed in the sensor node.

includes the signal preprocessing, feature extraction, and classification.

### 2.4.1 Preprocessing

Signal preprocessing involves removing components outside of the scope of a given evaluation as they can affect the feature extraction. For instance, High-pass/low-pass digital filters are typically applied to remove the DC offset and signal drift, high-frequency artifacts, and external noise in the form of PLI.

These can be implemented as Impulse Response (IIR) or Finite Impulse Response (FIR) filters. Typically IIR filters obtain a steeper attenuation than FIR filters while only requiring a few taps (or coefficients). Hence, reducing delay introduced on the final processing chain. Nonetheless, this filter cannot be used in some applications (such as ERP-based systems) as the phase response is not linear, and the filter stability cannot be guaranteed. FIR filters do not face these issues, but the magnitude response is shallower,

hence, requiring a large number of taps to achieve significant attenuation, which impacts the overall delay of the processing. Therefore, the filter implementation should take into consideration these issues to satisfy the application requirements.

Other preprocessing techniques include information reduction. Principal Component Analysis (PCA) is a widely used lossy statistical procedure to reduce the dimensionality of a set. PCA performs an orthogonal transformation to map the original data into a new set of values, called principal components. As the first components of this transformation have the most substantial variance (hence most of the information), only a few might be required to achieve the same performance. The channel reduction decreases the computational, therefore, it is a useful tool for embedded systems when the input data is significantly large [48].

The preprocessing also includes artifact removal techniques as artifacts might deteriorate the performance of the system, where techniques like regression and filtering and blind source separation[49] are typically used.

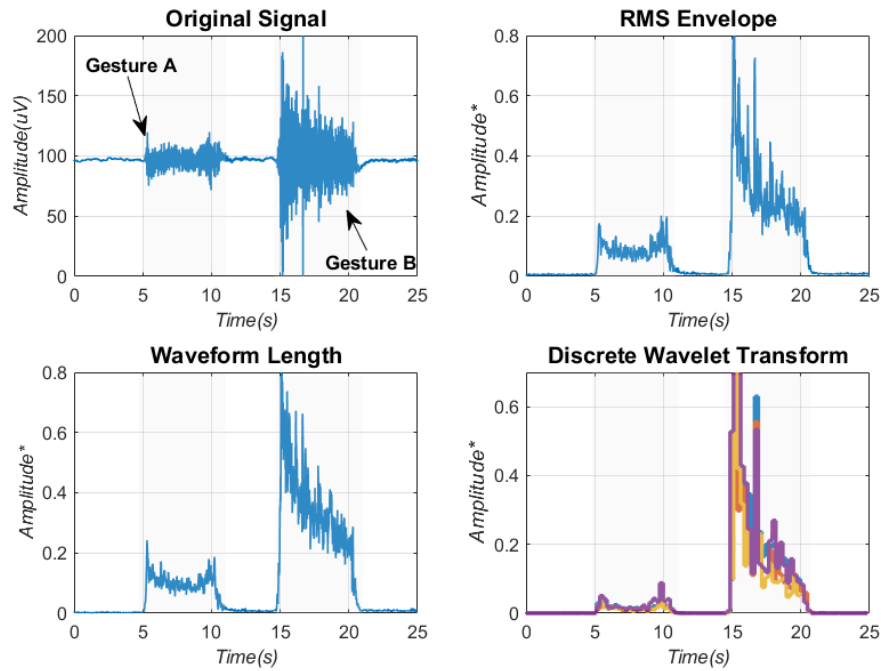
### 2.4.2 Feature Extraction

Feature extraction techniques allow transforming salient information contained in the signal into a condensed representation. This process is vital as it highlights the differences between different events, which allows for pattern recognition algorithms to provide significant performance. A feature extraction mechanism also handles the non-stationary behavior of biosignals by performing the feature transformation over a time window.

Biosignal feature extraction can be roughly divided into three different categories, namely, time, frequency, and time-frequency. Table 2.2 presents a summary of the

**Table 2.2:** Typical Feature Extraction Techniques for Biosignals

Signal	Time	Frequency	Time-Frequency
EMG[50, 51]	IntegratedEMG Mean Absolute Root-mean Square Waveform Length	Autoregrssive Coefficients Freq. Median Freq. Mean Modified Mean Frequency	Short-Time Fourier Transform Wavelet Transform Emp. Mode Decomposition
EEG[48, 52, 53]	Mean Standard Deviation Mean Absolute Value	PSD Autoregrssive Coefficients	Short-Time Fourier Transform CCA Wavelet transform
ECG[48, 52, 53]	QRS detection	PSD Autoregrssive Coefficients	Short-Time Fourier Transform Wavelet Transform



**Figure 2.9:** Example of the signal output after feature extraction of EMG signals. The top-left figure presents the original signal while the remaining present the different signal extraction outputs for RMS, Waveform Length, and Discrete Wavelet Transform.

preferred feature extraction techniques applied to EMG, ECG and EEG signals. Still, it is worth noticing that feature selection is not a deterministic task and obeys extensive analysis to minimize the total error of the complete system. Fig 2.9 shows examples of several features extracted from an EMG signal

### 2.4.3 Classification

Classification of biosignals involves inferring the state or condition of the human body based on the observed data. In this scope, the pattern recognition approach has demonstrated to be a robust way to overcome the limitations imposed by the non-linear behavior of the biosignal. Classification with pattern recognition methods is a well-studied field of research, with many contributions by machine learning communities.

A pattern recognition approach compares a previously defined model containing information about all mapped classes with an unseen new example. The class assignment of the new example relies on the proximity with the different data regions delimited during the training process.

Pattern recognition algorithms can be divided into two groups: supervised and unsupervised. Unsupervised algorithms attempt to learn patterns without prior information,



and it is used to differentiate unknown data structures. Clustering analysis, based on k-means and Gaussian mixture models, is widely used for this purpose. Supervised methods make use of prior information to identify the data regions. Typical supervised algorithms include logistic regression (LR), Support Vector Machines (SVM), and Neural Networks (NN). A logistic regression algorithm separates the class space linearly while SVM and NN are strongly non-linear.

SVM has gained popularity for embedded deployment as it provides superior classification performance to LR while being computationally lighter than NN based algorithms [40]. In this work, we have also explored the benefits of the High-Dimensional Computing, a brain-inspired approach that computes with points in the HD space (hypervectors) as an alternative to numbers, which offers SoA classification performance [54] while allowing for optimizations when coupled with the PULP architecture (bit-wise operations). Details of these algorithms and their implementation are introduced in Section 4.

## Chapter 3

# Hardware Design and Implementation of HMIs

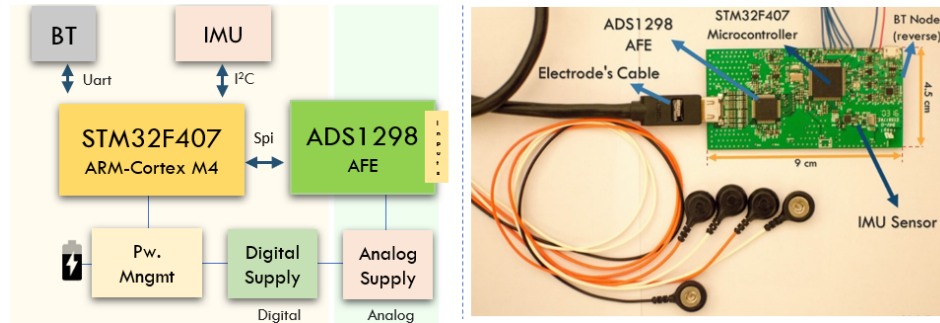
One of the most critical challenges in this field is to reduce the device form factor while performing all the computation online to provide unobtrusive monitoring. Current systems, nevertheless, are bulky, with limited portability. Similarly, the sensor interface is highly intrusive and does not provide adequate signal quality.

These issues are currently targeted by employing low-power non-intrusive wearable embedded design with flexible and dry sensor interfaces. Still, this task is not trivial since the deployment has to take into consideration resource constraints (memory and computational power) while also minimizing the power consumption to allow extended battery life.

In the following, we present two systems that tackle these issues. Both systems share the same AFE, namely, ADS1298, as it provides SoA performance while they differ in other aspects such as the processing unit and form factor. This chapter starts with the description of an ARM-based HMI and concludes with BioWolf, a PULP based HMI, which is the main focus of this thesis work.

### 3.1 ARM-based HMI

The first proposed IoT node for biosignal monitoring is based on the multichannel commercial Analog Front End (AFE) introduced in Section 2.2 (ADS1298)[55] connected with a low-power ARM Cortex M4 microcontroller. The board also includes other devices such as a barometric sensor and an inertial measurement unit (IMU). Fig. 3.1 shows a block diagram of the system (left) and the PCB implementation (right).



**Figure 3.1:** ARM BCI including Block diagram and PCB implementation.

### 3.1.1 PCB design

All the components are allocated in a 6-layers printed circuit board (PCB) with a single solid ground plane. To minimize current return paths, the power planes are split, keeping separated the analog and digital circuitry. Discrete components were carefully placed on both sides of the PCB to maximize signal integrity, maintaining a low level of noise and a small form-factor that results in 85x50 mm.

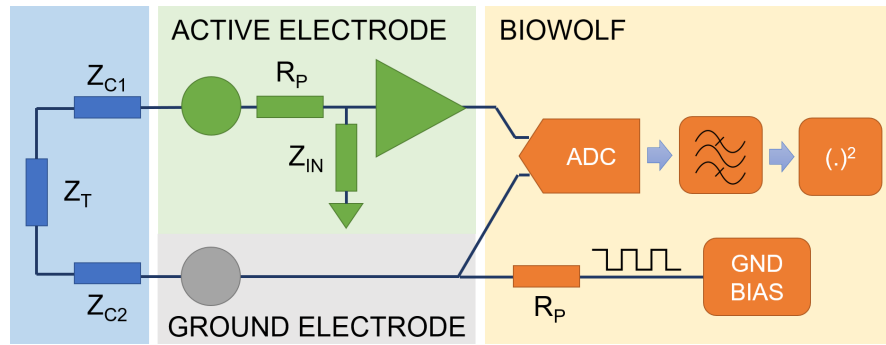
### 3.1.2 Sensor Interface

The eight channels of the ADC (described in Section 2.2) are connected with the active ExG sensors while the AFE's back-end streams the data via SPI to the microcontroller. This system, as the same as BioWolf, features zero-preparation dry electrodes, which are suitable for use outside clinical settings, due to the quick setup features, even for untrained users.

Several strategies are employed to tackle the typical challenges associated with dry electrodes (as introduced in Section 2.2). The first is placing the first amplification stage directly on the electrode to minimize the stray capacitance on the noise-sensitive electrode node, which reduces interference coupling through capacitive effects to the cables and artifacts due to cable movement while improving CMRR.

The reduction in the capacitive load on the electrode also improves system bandwidth as it reduces the voltage divider effect between electrode-tissue impedance and input capacitance itself, making it negligible even for dry electrodes showing very high values of contact impedance.

Since single-ended amplification of the electrode signal with a gain higher than one degrades the capability of the system to reject common-mode noise, only signal buffering is performed on the active electrode through a low-power, low-noise, rail-to-rail Operational Amplifier (O.A., namely AD8603 by Analog Devices) connected as a



**Figure 3.2:** Scheme of the electrode configuration used for both systems\*.

\* On BioWolf, a square wave is applied to patient ground to perform impedance estimation.

unity-gain buffer. Protection resistors ( $68\text{ K}\Omega$ ) limit the current delivered to the subject in case of a single fault condition (e.g., electrode shorted to positive or negative rail supply) to less than  $50\ \mu\text{A}$  as prescribed by applicable standards.


Another strategy, in this case, to minimize the number of wires, is to use the output signal of the amplification stage for providing power supply to the O.A. as well. This means that the positive power supply of the O.A. is at the same voltage of the inverting and non-inverting inputs, hence the need for a rail-to-rail operational amplifier. A forward-biased diode connects the output pin of the O.A. to the power supply. The output is biased by the following stage, which entails a  $10\text{ k}\Omega$  resistor toward a  $2.7\text{ V}$  power supply. The output of the amplification stage can swing between  $1.8\text{ V}$ , which is the minimum power supply voltage of the O.A. and  $2.2\text{ V}$ , which corresponds to the  $50\ \mu\text{A}$  quiescent current required by the O.A. Fig.3.2 presents a block diagram of the complete sensor interface.

### 3.1.3 Subsystems

The board also has an IC for power management for automatic power source switching (battery or USB). Analog, digital, and communication subsystems are supplied by separate low-dropout voltage regulators, which enables finer power management strategies, like duty-cycling submodules as to enhance battery life. Finally, data can be streamed to an external device through a BT 2.1 module. Table 3.1 summarizes all the features of the system.

## 3.2 BioWolf: The Parallel Ultra Low Power HMI

BioWolf is an integrated platform for computationally-intensive medical IoT applications, which addresses the typical challenges for biosignal processing by providing a



**Figure 3.3:** BioWolf System. Block diagram (left) and PCB implementation (left).

ULP compute platform that can process biosignals in parallel and locally with a power budget lower than that required by the analog front-end (AFE) to acquire the data. The platform is based on Mr. Wolf [56] and a commercial Bluetooth Low Energy (BLE) SoC (Nordic nRF52832) that enables data transmission and auxiliary support for the system, an 8-channel Analog Front End (AFE) to convert the EEG signals (ADS1298) and an Energy Harvesting (EH) subsystem for charging battery from solar source.

### 3.2.1 PCB design

All the components are assembled in a 20x40 mm form factored 4-layer Printed Circuit Board (PCB) that aims to provide full portability and wearability. Fig. 3.3 shows a block diagram of the system and the PCB implementation, while Fig. 3.4 shows BioWolf encased in a patch-like form factor and worn on the head in an experiment of SSVEP-based BCI for hands-free control of a digital score for musicians.

<b>AFE (ADS1298)</b>	
Channels	8
Input reference Noise	4(micro)Vpp
Power Consumption	12.5 mW
No. of adjustable gains	7
Signal-to-noise Ratio	112dB
Resolution	24bit
<b>IMU (MPU-9150)</b>	
DOF	9
Resolution	16bit
Sample Rate(Acc)	1Khz
Power Consumption	1.4mW
<b>Microcontroller (STM32F407vgt6)</b>	
Operational Frequency	168Mhz
RAM	192kB
Flash Memory	1MB
Run mode Pw. Consumption	152mW
Sleep mode Pw. Consumption	39mW

**Table 3.1:** Relevant features of the devices used in the system.



**Figure 3.4:** Example of BioWolf montage in an SSVEP-based BCI for hands-free control of a digital score for musicians.

### 3.2.2 Sensor Interface and electrical characterization

As mentioned before, BioWolf shares the same sensor interface as the ARM-based HMI (ADS1298 and zero preparation electrodes). We characterized the complete system with the AFE running at 500 samples-per-second (SPS), which provides a -3 dB cutoff frequency of 131 Hz, exceeding the needs of most applications. As the active electrode bandwidth exceeds 100 kHz, this value is essentially determined by the AFE internal filters. Noise is measured by shorting the inputs of the electrodes and varies depending on the chosen PGA gain.

We consider three different configurations: without active electrodes and with active electrodes in single-ended (negative input of the ADC connected to patient ground) and differential signal acquisition (positive and negative input connected to two separate active electrodes).

In the first configuration, with PGA gain equal to 1, noise is measured at  $1.62 \mu V_{RMS}$  in the 0.5-100 Hz band, decreasing to  $0.95 \mu V_{RMS}$  (gain = 2),  $0.48 \mu V_{RMS}$  (gain = 4) and  $0.40 \mu V_{RMS}$  (gain = 12). Common Mode Rejection Ratio for a 50 Hz signal ranges from a minimum of 115 dB (gain = 1) to 122 dB (gain = 12).

It is worth mention that for the second configuration, i.e., using active electrodes, there is a slight degradation in noise and CMRR performances due to the additional noise introduced by the electrode circuitry and the finite GBW of the O.A., which results in each electrode having a gain slightly different from the theoretical value of 1.

As the input swing with two-wires active electrodes is limited to a few hundred mV, we characterize the system only for a gain of 6 and 12. In single-ended configuration,

noise increases at  $0.63 \mu V_{RMS}$  for gain 6 and  $0.6 \mu V_{RMS}$  for gain 12. In the differential configuration, due to the presence of a second O.A., noise increases to  $0.77 \mu V_{RMS}$  and  $0.75 \mu V_{RMS}$  respectively. CMRR is measured at at 92 dB (gain = 6) and 96 dB (gain = 12). Channel isolation exceeds 100 dB in all cases.

These values are in line with IFCN standards for clinical recording of EEG signals [32], which are generally considered as the most stringent for bio-potential acquisition, demonstrating how ADS1298 is a feasible choice for medical-grade signal acquisition.

The ADS1298 can perform lead-off detection by injecting a small current between two electrodes and measuring the resulting voltage difference. This is a typical approach in systems based on passive electrodes, or active electrodes with custom ICs [34] but is not a viable option in systems employing active electrodes based on commercial OA, where current sources should be added at the input of each active electrode amplifier, increasing the cost and size of each electrode.

To overcome this limitation, we apply the scheme depicted in Fig. 3.2, where  $Z_{C1}$  and  $Z_{C2}$  represent the contact impedance of the measuring and ground electrodes, respectively.  $Z_T$  is the equivalent impedance of the tissue between the two electrodes and is negligible as being significantly smaller than dry-electrode contact impedance.  $R_P$  represents the protection resistors added on each electrode, intended to limit the current to the subject in case of a single fault, which is considered as an add-up to contact impedance.

When contact impedance measurement is required, the system superimposes a 200 mV square wave to the DC voltage on the ground electrode. As the amplifier is extremely close to the electrode, parasitic capacitance on the electrode are almost entirely due to the input capacitance of the amplifier itself ( $Z_{IN}$ ) and, as the values of  $Z_{C1}$ ,  $Z_{C2}$  increase, it causes a voltage divider effect which reduces the amplitude of the first harmonic of the square wave. This voltage divider effect is directly responsible for most of the common-mode interference and can provide an indirect measure of the contact impedance. The difference between the signal's full amplitude on the ground electrode and that detected on the signal electrode ultimately provides a measure of the contact quality and can be computed through the average power of the first harmonic (signal is band-pass filtered). The higher the value, the higher the contact impedance and, therefore, the worse the contact quality. This information can be compared to a threshold (which can depend on application and type of electrode) to inform the user whether the quality is good enough for the BCI to work as expected.

We verified the behavior of the contact impedance measuring circuit by connecting impedances of known value between the ground electrode and one active electrode.

Contact impedances  $Z_{C1}$  and  $Z_{C2}$  (Fig.3.2) are modeled as the parallel of a capacitance and a resistor, which is a typical model to represent an interface between an electrode and skin [33].

Values have been chosen to be those typical of wet electrodes with skin preparation (47 nF || 51 K $\Omega$ ), and representing worse quality contact such as that typical of dry electrode setups (4.7 nF || 510 K $\Omega$  and 1 nF || 2.2 M $\Omega$ ). Input impedance  $Z_{IN}$  is represented by a capacitance of 20 pF or 120 pF. The first value represents a reasonable value of input capacitance for an active electrode, the second one for a standard setup with approximately 1 m cables connecting electrodes to amplification circuits. Table 3.2 compares the theoretical and measured amplitude of the 40 Hz harmonic at the ADC input due to the voltage divider effect between contact impedance  $Z_{C1} + Z_{C2}$  and impedance to ground  $Z_{IN}$ . Signal amplitude is expressed as a fraction of the input signal. Results show excellent agreement between theoretical and measured values, validating our approach.

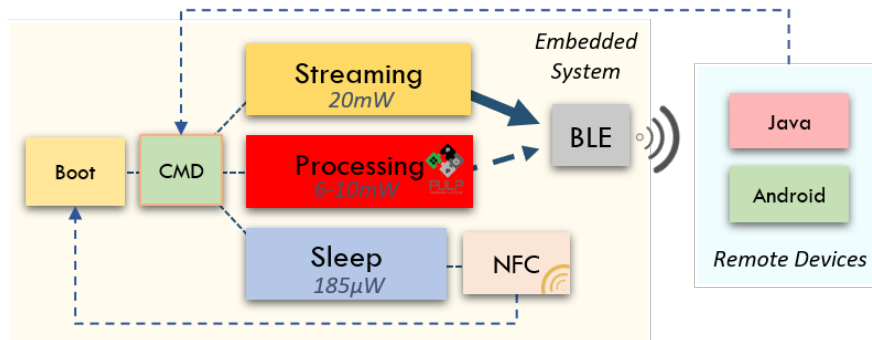
### 3.2.3 Operational Modes

Thanks to the presence of both a standard MCU/SoC and Mr. Wolf Ultra-Low Power (ULP) processor, BioWolf behaves as a highly configurable platform that can work in three different modes, i.e., *Stream*, *Processing* and *Deep-Sleep*. In stream mode, Mr. Wolf goes into sleep mode, and the Nordic SoC acts as the Master of the SPI bus to read data directly from the ADC. The samples can then be streamed using the BLE link embedded in the SoC. This mode is useful for BCI algorithm development and testing, where raw data is required for offline data analysis. In processing mode, Mr. Wolf ensures the best system-level energy efficiency as computationally intensive processing is required and is therefore adopted in BCI applications. Mr. Wolf acts as the Master on the SPI bus and directly samples data from the ADC. The processing required by the BCI application is then executed on this ULP processor, while the other devices act as slaves. Results from the processing can be transmitted using the BLE link through the

**Table 3.2:** Contact impedance measurement for different values of  $Z_{C1,2}$  and  $Z_{IN}$

$Z_{C1,2}$	$Z_{IN}$	Expected	Measured	Estimated $ Z_{C1,2} $
47 nF    51 K $\Omega$	20 pF	0.21	0.21	43.9 K $\Omega$
47 nF    51 K $\Omega$	120 pF	1.21	1.20	43.4 K $\Omega$
4.7 nF    510 K $\Omega$	20 pF	2.0	2.0	83.3 K $\Omega$
4.7 nF    510 K $\Omega$	120 pF	1.20 %	1.21 %	83.9 K $\Omega$
1 nF    2.2 M $\Omega$	20 pF	0.93 %	0.94 %	1.92 M $\Omega$
1 nF    2.2 M $\Omega$	120 pF	5.33 %	5.76 %	1.99 M $\Omega$





**Figure 3.5:** Operational modes of BioWolf.

Nordic SoC. The complete system can be put into a deep sleep mode to minimize power consumption while the system is in standby. Fig. 3.5 summarizes the functionalities and operation modes introduced above.


### 3.2.4 Subsystems

The nRF52832 SoC from Nordic provides data communication to the system. The Arm Cortex-M4 MCU (64 MHz) allows flexible Bluetooth 5 (BLE) communication at a low-power budget. Power consumption for the radio is 15.9 mW when streaming data at 1 Mbps, 0 dBm output power, and 16.2 mW when receiving data at the same rate. Its ULP features allow reducing power consumption down to  $4.5 \mu\text{W}$  with the system on full RAM retention, wake on any event (such as interrupts from the AFE when new data is available). The power consumption can be further reduced down to  $2.1 \mu\text{W}$  when the complete system is in sleep mode.

The nRF52832 SoC also serves as a device manager of the board, allowing to choose the operation mode (sleep, stream, and processing), and to program Mr. Wolf by directly accessing its volatile memory through a JTAG interface. The SoC also monitors the battery status through a fuel gauge IC. Additionally, the SoC allows for Near-Field Communication (NFC) that is used to wake-up the complete system after entering sleep mode by tapping the board with an NFC-enabled device.

### 3.2.5 Energy Harvesting

A Texas Instruments BQ25570 manages power supply, battery charging, and energy harvesting. The IC implements Maximum Power Point Tracking (MPPT) to maximize energy conversion efficiency for all lighting conditions, up to 90%. The harvested energy is used to recharge a 65 mAh battery with a size of 29x12x2 mm. The Energy Harvesting Subsystem (EH) also provides a highly efficient buck converter that outputs the 1.8 V



**Figure 3.6:** Solar Panel current charging output for different illumination conditions. Indoor illumination is typically around 600 lux (magnified), while in outdoors, the illumination is about 10k lux.

required to supply the digital portions of the board. An additional output voltage of 3 V is available, and it is used to power an LDO, which supplies a 2.7 V rail to the AFE and the remaining analog portions of the system.

In indoor conditions (office illumination), the complete EH subsystem can provide up to 0.4 mW of power, enough to sustain the device during deep-sleep. In outdoor environments, the solar panels can generate up to 5 mW. Thus, under ideal conditions, the EH subsystem is capable of extending the battery duration of the application or of providing self-sustainability, as it is the case of some of the applications presented in the following. Fig. 3.6 shows the measured current generated for different illumination conditions.

### 3.2.6 Software

Fig. 3.7 shows the GUI that provides support for the final user. The software, running on JavaVM, allows visualizing and storing the incoming data. The software also allows changing the configuration of the ADS1298 on-the-fly (op. mode, sample rate, and gain) and visualization parameters (data scale). The system supports real-time filtering (bandpass and notch) to remove artifacts or to visualize, for instance, the EEG rhythms such as Alpha.

Other add-ons include battery monitoring and trigger support. The later can be generated manually (button) and remotely through a TCP/IP server. The system also allows visualizing the outputs of Mr. Wolf when the board operates in Processing mode through a serial terminal. The BLE connection from the board to the computer is

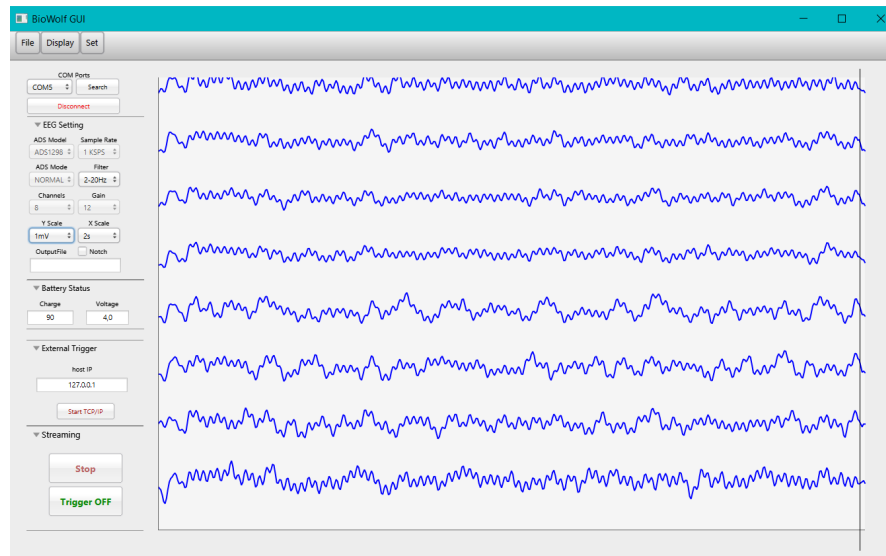


Figure 3.7: BioWolf GUI for streaming and data storing (PC).




Figure 3.8: BioWolf GUI for Android during EMG classification.

achieved through a USB dongle. An Android application has also been developed to provide more flexibility and freedom of movement. Fig. 3.8 shows a screenshot of the running application that shares the same features introduced for the PC version.

### 3.2.7 Discussion

In this chapter, two HMI systems have been presented. BioWolf offers significant advantages to the ARM-based HMI and the SoA devices for onboard biomedical processing, which are more evident when comparing it with the previous work presented in Table 3.3. Although some of the listed systems can provide higher performance ([MFlop/s]), they lack efficiency ([MFlop/s/mW]), where BioWolf outperforms all systems by a factor of 10.

**Table 3.3:** Comparison between SoA embedded platforms for BCI processing.

	Application	Footprint [mm x mm]	Number of Channels	Connectivity	System Power [mW]	System On Chip	Performance [MFlop/s] <sup>1</sup>	Efficiency [MFlop/s/mW]
Condori [59] :	Motor Imagery	83x59	3	n.a.	>10000	Exynos 5422	13600	<1.36 <sup>2</sup>
Lin [60] :	SSVEP	n.a.	2	BT	5000 <sup>2</sup>	Qualcomm Snapdragon 800	9200	1.8 <sup>2</sup>
Li [61] :	$\alpha$ -band power/SVM	37x56	2	Wi-Fi/BLE	321.6	Qualcomm Snapdragon 400	4800	1.8 <sup>2</sup>
De Venuto [12] :	P300	85x56 <sup>3</sup>	32	Wi-Fi	>1200	BCM2836	3600	3 <sup>2</sup>
Chi [62] :	SSVEP	n.a.	8	BT	1000	TI OMAP 2420	434	0.5 <sup>2</sup>
Chai [57] :	SSVEP	36x36 <sup>3</sup>	2	2.4GHz RF	>3300	Nordic nRF24	16 <sup>4</sup>	1 <sup>4</sup>
McCrimmon [63] :	Motor Imagery	130x90	8	n.a.	1000	Atmel SAM3X8E	84 <sup>4</sup>	1 <sup>4</sup>
Salvaro [41] :	SSVEP	85x50	8	BT2.1	27.50	ARM-STM32F407	168	1.2
Kartsch [58] :	Drowsiness	85x50	8	BT2.1	109.00	ARM-STM32F407	168	1.2
<b>This work :</b>	<b>SSVEP</b>	<b>40x20</b>	<b>3<sup>5</sup></b>	<b>BLE</b>	<b>6.31</b>	<b>Mr. Wolf</b>	<b>1400</b> <sup>6</sup>	<b>18<sup>7</sup></b>

<sup>1</sup> Ideal (i.e. considering maximum operating frequency x # of FPUs).<sup>2</sup> Estimated.<sup>3</sup> The complete system is composed of multiple embedded devices.<sup>4</sup> Fixed-point.<sup>5</sup> 8-Channel Capable.<sup>6</sup> Considering two floating-point operations for fused multiply and accumulate. Cluster 1.1V, 350 MHz.<sup>7</sup> Operating point: 0.8V, 100 MHz.

A significant advantage of BioWolf is the reduction of intrusiveness. Although [57] offers a similar form factor, this device only serves as streaming. Hence, under the same conditions, only the ARM-based BCI (also in [41, 58]) presented in this section is comparable, where BioWolf offers a 4.25x smaller area.

Another notable benefit of BioWolf is its high versatility, allowing for both low-complexity and computationally intensive applications to run efficiently through single-core or multi-core processing.

To complement the advantages in terms of energy efficiency (and therefore also size, weight, and wearability of the system), we have also demonstrated a subsystem for energy harvesting from solar sources and one for checking electrode contact quality, both for passive and active electrodes. The compactness and ease of use of our platform allow it to be seen as a valid self-contained wearable device for day-to-day BCI applications.

Details about the power consumption are introduced in the following sections as they depend on the specific application.

## Chapter 4

# HMI applications using BioWolf

The following chapter presents four case studies for HMI. The first three (EEG) include a drowsiness detection system, a CCA-based BCI for SSVEPs, and a system for food quality grading through ERPs, while the fourth introduces an EMG gesture recognition system based on High-Dimensional Computing. Details about the experimentation, offline data analysis, embedded deployment, and optimizations are introduced, demonstrating the efforts to produce a system that offers both unobtrusiveness and energy efficiency. Each section also offers a discussion to highlight the improvements achieved for each application by employing the resources available in BioWolf.

### 4.1 Drowsiness Detection

Drowsiness detection mechanisms gained attention recently as it is one of the prevalent causes of accidents within the mining, driving, and industrial activities. Drowsiness is typically quantified using behavioral analyses based on camera eye-tracking systems as well as analyzing physiological features like ECG. EEG signals have also reached more notoriety recently.

Detection systems typically use specific drowsiness indicators from only one of these methods, leaving a risk of missed detection since not all the population presents some symptoms of drowsiness [64]. Hence, multi-feature systems have the potential to provide a more robust detection. Nevertheless, this approach comes with a more significant system complexity, especially when implemented on the current power-hungry SoA platforms, which only have a slim chance of providing useful battery lifetime and wearability.

The following work presents a drowsiness detection scheme fusing behavioral information coming from user motion through an IMU sensor and physiological information

coming from brain activity through a single EEG electrode. The system is implemented and tested on a low-power programmable platform based on an ARM Cortex-M4 microcontroller (introduced in Section 3.1). It is capable of detecting 5 different levels of drowsiness with an average accuracy of 95.2% and a battery life of 6 hours while offering a more robust system (through multiple drowsiness features) in a highly embedded platform.

In this first part, the energy optimizations achievable by accelerating the sensor fusion-based drowsiness detector on a parallel ultra-low power (PULP) platform have also been studied, demonstrating that this architectural shift can lead to improvements (>7x) in battery life to SoA systems.

The section concludes with the implementation of a modified version of the original drowsiness detection system on BioWolf.

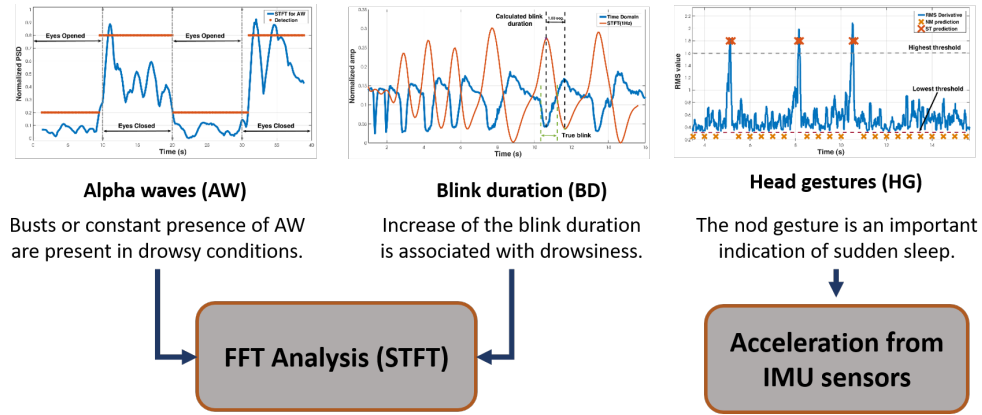
#### 4.1.1 Introduction

Detection of driver's drowsiness is an active research field in both industrial and transportation areas. Several automotive companies, universities, research centers, and governments are contributing to the development of Advanced Driver Assistance Systems (ADAS), which aims to analyze the different technologies and techniques to reduce the risk of accidents caused by drowsiness [65, 66].

Commonly, commercial vehicles identify safety risks by analyzing the driver's behavior, which includes the monitoring of the vehicle's position to the lane markings [67]. Nevertheless, these systems fail to alert the driver before the occurrence of an incident. A more reliable approach is to measure physical behaviors (PERCLOS, blink frequency, nodding, and yawing occurrences) to detect the level of fatigue.

The work presented by Flores et al. [68] relies on a digital camera to compute the driver's eyes state. This architecture is coupled with NIR (near infra-red) illumination and stereo vision to deal with low-light conditions [69, 70]. Still, computer vision techniques result in a challenging task due to the variability of environmental factors and the computational requirements needed to provide a real-time assessment.

A new trend in research is to directly measure biometric signals such as EEG, EOG (Electrooculography), ECG, PPG (Photoplethysmogram), and eventually fusing data from multiple sensors [71, 72] to improve the robustness [73–75]. The system proposed by Lee et al. [76] relies on sensor fusion algorithms to detect the driver's fatigue level using ECG, heart rate variability, blood pressure, and PPG signals. Reyes et al. [77], proposed the integration of body area sensors and vehicular ad-hoc networks for traffic safety using



**Figure 4.1:** Features extracted from EEG and IMU sensors.

a wireless physiological signal-acquisition module and onboard PC processing. Similarly, Lin et al. [78] proposed a system based on EEG combining independent component analysis (ICA), power-spectrum analysis, correlation evaluations, and linear regression model. Notwithstanding, the approaches above rely on cumbersome high-end computing platforms, which discourages its use.

Another non-exploited advantage of physiological monitoring is the lower computational complexity required (to computer vision methods) that allows for the use of embedded microcontrollers[79–81].

This work, taking into account the previous research, exploits a sensor fusion algorithm based on IMU and EEG sensors integrated on a single wearable board to provide a 5-level drowsiness alert system, covering a wide variety of drowsiness-related symptoms.

The following introduces the drowsiness parameters extracted from both EEG and IMU and the techniques to fuse the information and the onboard/online implementation of the system. Furthermore, the detection algorithm has been implemented on a near-threshold parallel platform to evaluate possible energy savings to previously reported architectures.

#### 4.1.2 Feature Extraction

In this application, the blink duration and alpha wave activity are used and extracted from EEG, as they correlate with the presence of drowsiness [82–86]. The head posture also correlates with drowsiness[87–89], and it is extracted through an IMU sensor.

Blink duration is estimated by analyzing the lower spectral components of the EEG through the Short-Time Fast Fourier Transform (STFT) (over 512 samples with 32-samples overlap). Fig. 4.1 (center) shows the energy of the STFT 1Hz during a series of

eye blinks. The presence of a high and low peak on the signal is used for the estimation as they correlate with the actual blinks. The amplitude difference compared to a threshold helps in removing artifacts such as short blinks, eye movements, and saccades (No blink, NB). Table 4.1 shows the results of the validation of this method for short (SB), middle (MB) and, long blinks (LB).

The alpha wave activity is estimated using the previously calculated STFT. An example of the alpha power during the closure of the eyes is presented in Fig. 4.1 (left). Following the Objective Sleepiness Scoring (OSS) [90], it is possible to identify a drowsy state, without referring to more evident behaviors, such as the closure of the eyes, by evaluating the period on which the alpha waves level exceeds a given threshold, over a time window of 20 seconds. In this work, we refer to this feature as Cumulative alpha wave (CAW). This feature triggers a binary indicator, which is true when alpha waves activity is at least present for 5 seconds on the 20-second time window specified by the OSS. This is achieved by accumulating the value of the alpha wave occurrences above threshold (in seconds). The same approach is used to detect the closure of the eyes but using instead a shorter time window (3 seconds), named Continuous Alpha Wave (CoAW), in this work. In this case, the alpha activity must remain above the threshold during the complete window to be considered. The threshold for both features was estimated empirically by comparing alpha waves activity during regular morning activity and drowsy conditions, i.e., during nightly experimentation.

The nodding gesture is detected through the RMS of the derivative of the accelerometer signal on the three dimensions (32 samples, 1 sample overlap) to assert three different states: sudden tilt (ST), normal movement (NM), and no movement (NOM). Two thresholds are compared to the RMS value to assign one of the three activity states. If the RMS is greater than the higher threshold, the ST state is asserted. If the RMS is in between the two thresholds, the NM state is asserted. If the RMS is below the lower threshold, the NOM state is asserted. The thresholds have been selected empirically. The accuracy of the classifier for the three different classes is reported in Table 4.2. Fig. 4.1 (left) show an example of the extracted features from the IMU sensor, the thresholds selected, and the detection output.

### 4.1.3 Sensor Fusion

The sensor fusion algorithms input the extracted values to obtain the drowsiness level classification using several thresholds. Fig. 4.2 summarizes the complete sensor fusion algorithm. The different levels are asserted when the following events/conditions are present:



- **Level 1: Normal movement**

At the first level, only the IMU sensor is used. The system remains in this level when the blink duration is below 500ms, the user produces no alpha waves (or below threshold), and the IMU classifier reports No movement (NOM) or Normal movement (NM).

- **Level 2: Blink Duration + No movement**

Two conditions are required to trigger the alarm at this level. The blink duration must be longer than 500 ms, according to previous investigations [82–84] and the gesture classification must return no movement (NM) class, used to confirm the reduced activity triggered by a drowsiness state of the user. If normal movement (NM) is reported, the value of the blink duration is set to zero, assuming that the user is not to be under drowsy conditions.

- **Level 3: Alpha waves burst detection**

The occurrence of alpha waves bursts increases when the user enters in a deeper drowsiness state. Following the Objective Sleepiness Scoring (OSS) [90], it is possible to identify a drowsy state by evaluating the period on which the alpha waves level exceeds a given threshold, over a time window of 20 seconds. The alpha waves are quantified using the PSD method described before, and the alarm is triggered if the value is above the threshold for at least 5 seconds on the time window.

- **Level 4: Sudden tilt (nodding)**

**Table 4.1:** Normalized peak-to-peak difference for different blinks and detection accuracy.

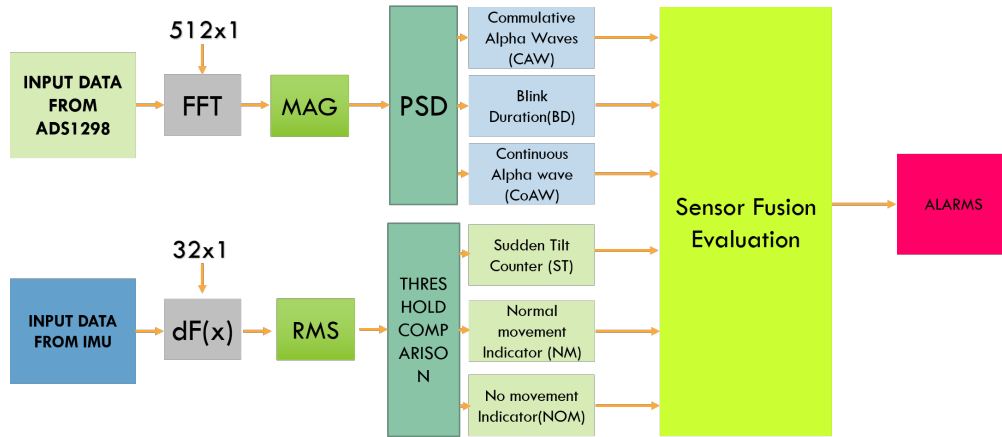
Trial No.	SB <sup>+</sup>	MB <sup>+</sup>	LB <sup>+</sup>	NB <sup>+</sup>	Acc.* %
Trial 1	0.03	0.168	0.308	0.01	92
Trial 2	0.03	0.168	0.290	0.009	93
Trial 3	0.03	0.167	0.302	0.01	87
Avg. time	0.367ms	0.61ms	1.174ms	-	-

\* The accuracy was evaluated comparing the estimated values with the values obtained using the definition of blink given by [91].  
+ normalized values


**Table 4.2:** Threshold values for ST, NOM and NM calculated from the IMU sensor.

Gesture	Threshold	Accuracy
ST	Output >1.6	93
NM	0.4 >= Output <1.8	95
NOM	Output <0.3	97

The values here exposed correspond to the threshold giving the highest accuracy values.



**Figure 4.2:** Block diagram of the extracted features to compose the drowsiness sensor fusion algorithm.



**Figure 4.3:** Block diagram of the complete wearable deployment for drowsiness detection based on the ARM Cortex-M4 MCU.

Level 4 of the drowsiness state is asserted using only the 3-axis accelerometer to recognize nodding movements of the user's head. In this alarm level, only the sudden tilt is accounted. A single event is needed to activate the alarm, since the nodding gesture represents a stronger indication of drowsiness [87–89].

- **Level 5: Constant presence of alpha waves**

A constant burst of alpha waves is an indication of the closure of the eyes [92] induced by a total loss of attention or sudden sleep. Such a state is detected using the same feature extraction method reported for level 3. More specifically, if the maximum energy of the alpha waves stays over the threshold for more than 3 seconds, the alarm level 5 is asserted. This was considered as the highest level of alarm since it corresponds to the most dangerous situation where the user is falling asleep.

#### 4.1.4 Embedded Deployment

The complete processing and classification algorithm is implemented on the ARM-based architecture presented in Section 3.1. Once the signals have been acquired, the




Figure 4.4: Detection accuracy of the different drowsiness alarms.

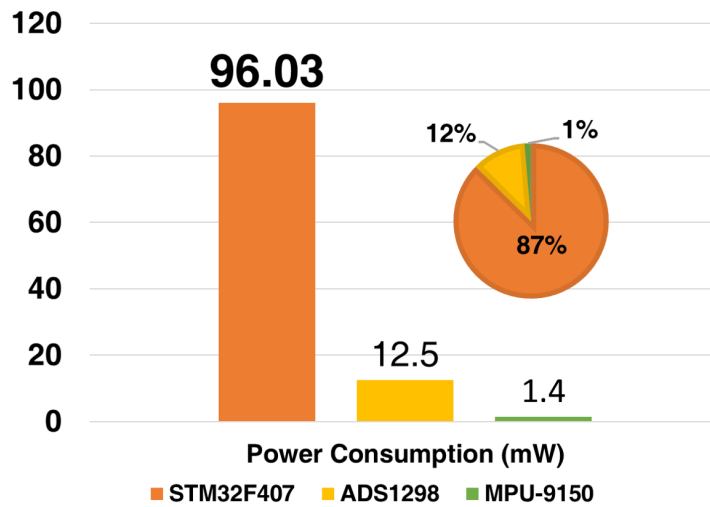
power spectral density (PSD) of the EEG signal and the RMS of the IMU signal are calculated by exploiting the optimized Cortex microcontroller Software Interface Standard (CMSIS) DSP Software Library [93]. Detected drowsiness levels can be streamed via the BT module to an Android application, developed to show the system state and the intermediate or final outputs of the processing chain. The sensor interface is based on auto adhesive circular gel-based electrodes, and the board is secured on the head using rubber bands. Fig. 4.3 shows a high-level diagram of the complete wearable deployment.

#### 4.1.5 Experimental Results

Testing a drowsiness detection system is not a trivial task. A real scenario where a drowsy person drives a car exposes the test subject to dangerous situations. Hence, validation of the proposed system is performed simulating the behavioral and physiological drowsiness symptoms, rather than in real conditions. This approach has already been validated by previous works and demonstrated to be an effective and reliable way to test drowsiness [94–97].

The 5-level alarm system described previously is tested on ten healthy subjects with no previous history of neural diseases (aged  $32 \pm 5$  years old). The participants were under sleep deprivation (3 hours of sleep) the day before the test. The tests were conducted at late night hours to maximize the drowsiness effects.

The board is placed on the subject’s head. The electrodes are located at Oz (Positive electrode) and Fpz (Negative electrode), following the 10-20 reference system [98], with a reference electrode placed in A1. Locations are chosen to maximize SNR in the



**Figure 4.5:** Power consumption of the embedded implementation (on PCB). The MCU dominates the power consumption, employing 87% of the total power.


acquisition of both relevant brain activity rhythms and eye blinking as needed by the application and to be compatible with integration on head caps, helmets and headbands. Subjects reported that the device did not cause significant discomfort after 30min of testing.

During the test, each alarm level is assessed separately, leading to a binary evaluation of each level (fail/pass). Each test subject simulates the alarm conditions five times. The first level is identified as fail if any alarm is triggered over 1 minute of regular activity performed by the user. The second level is asserted as a fails if the alarm is not triggered after closing the eyes for more than 500ms while performing no movement. Regarding the third level, the test fails if the alarm is not triggered after the corresponding evaluation window (20 seconds). In the case of the fourth level, the detection must be done on a single nodding event. For the last level, the detection fails if the system does not detect the condition after three seconds from the closure of the eyes.

Fig. 4.4 shows the experimental results, where average accuracy reaches to 95.2%. It is noteworthy that the lowest accuracy is measured in pure alpha wave detection, confirming the high variability of purely physiological detection and the added value of a sensor fusion approach to improve the robustness of detection.

#### 4.1.6 Embedded System Performance

Fig. 4.5 shows the system power consumption, including the contributions of the ADC, IMU, the microcontroller (ARM-M4 @168Mhz), and the BT radio. When possible, the microcontroller can be put in sleep mode to reduce the average power consumption. Fig. 4.6 shows detailed information about the current drawn by MCU throughout the



**Figure 4.6:** Current drawn by the MCU during the processing and classification of the drowsiness levels.

processing and classification of the signal, denoting in fine detail power management applied to reduce power consumption. Thanks to this, the system reaches up to 6 hours of life with a 200mAh Li-Ion battery.

#### 4.1.7 Implementation on the PULP architecture

The kernel functions have also been implemented in a PULP architecture (PulpV3, described in Section 2.3.2) to evaluate the possible benefits (energy efficiency) when switching architectures.

Since the most compute demanding part of the algorithm is the FFT, a great effort has been spent to optimize this kernel on the PULP architecture, exploiting a fine-grained data-parallel scheme supported by the programming model. The FFT algorithm requires to compute a set of butterflies (i.e., partial FFT transforms) on the  $N$  input samples (where  $N$  is the size of FFT, 512 for this application) for each stage of computation, and the number of stages is equal to the 2-base logarithm of the number of input samples (9 in this case). In the baseline radix-2 algorithm, after each stage,

**Table 4.3:** Number of cycles required to compute each function in the analyzed embedded computing platforms [KCycles].

Kernel Func	ARM	Pulp SC	4CPulp	ARM/PulpSC	ARM/4CPulp	PulpSC/4CPulp
FFT RDX8	42.90	64.03	16.69	0.67	2.57	3.84
Magnitude	17.87	15.17	3.83	1.18	4.66	3.96
RMS (32s)	0.26	0.30	0.16	0.86	1.67	1.94
RMS(512s)	2.52	2.95	0.83	0.86	3.05	3.56

The FFT calculation is the most computationally demanding operation. Given its highly parallelizable computational algorithm, the 4-core PULP MCU reaches nearly ideal speedup.


data generated by the butterflies of the previous stage has to be shuffled to compute the butterflies on the following stage.

On the PULP architecture, the FFT is computed by splitting the butterflies calculation homogeneously among the four cores, and synchronizing the cores using hardware barriers after each stage of butterflies to maintain the data consistency. Two different implementations have been evaluated, described in the following.

As explained above, the baseline radix-2 FFT requires a synchronization barrier after each stage of butterflies (i.e. 9 barriers), leading to a relevant synchronization overhead not amortized by the small computational load required by each stage of butterflies. A more optimized approach relies on the radix-8 algorithm. Exploiting this implementation, each butterfly performs a single Discrete Fourier Transform (DFT) among 8 samples instead of 2 as in the case of the radix-2 implementation. This reduces the number of butterflies to be computed at each stage but it increases the computational complexity of each butterfly. The proposed implementation is composed of 3 stages each with 64 butterflies (16 for each core), therefore 3 barriers are triggered to accomplish the full 512 samples FFT. Hence, this approach increases the available parallelism and reduces the synchronization overhead with respect to the radix-2 algorithm. The computation of the magnitude of the FFT is parallelized by dividing the signal by 4, thus each core works independently without synchronizations. Both the FFT and magnitude parallel implementations feature a speedup greater than 3.9 w.r.t. the single core version on PULP, showing a quasi-ideal parallel speedup in performance.

Regarding the computation part related to the IMU signal, the RMS envelope is computed on the last 512 samples instead to 32 since it offers more efficient use of cores. A parallel version was implemented by splitting again the signal in four parts. Each core computes the summation of the square of the signal assigned and finally the cores are synchronized. In the last phase, a single core is in charge to sum the four results and to compute the square root.

Table 4.4 offers a comparison between different platforms for real-time operation, where core frequency is modulated to achieve real-time. This table includes a high-end MCU STM32F407x (original MCU) and an ultra-low-power MCU Ambiq Apollo, both based on Cortex-M4 processor, and PULP executing on a single core and four cores. Although Ambiq Apollo offers superior energy efficiency, it cannot satisfy the real-time constraints due to its limited maximum operating frequency (24 MHz). On PULP, the frequency required to maintain the latency obviously decreases when increasing the number of cores. This, coupled with the near-threshold computing capabilities of the PULP platform aim towards a significant improvement in energy efficiency. This concept is well highlighted in Fig. 4.7, which also shows the comparison with off-the-shelf MCUs



**Figure 4.7:** Power consumption comparison between different platforms. The 4-cores PULP MCU offers 63x energy saving with respect to the implementation on STM32F407.

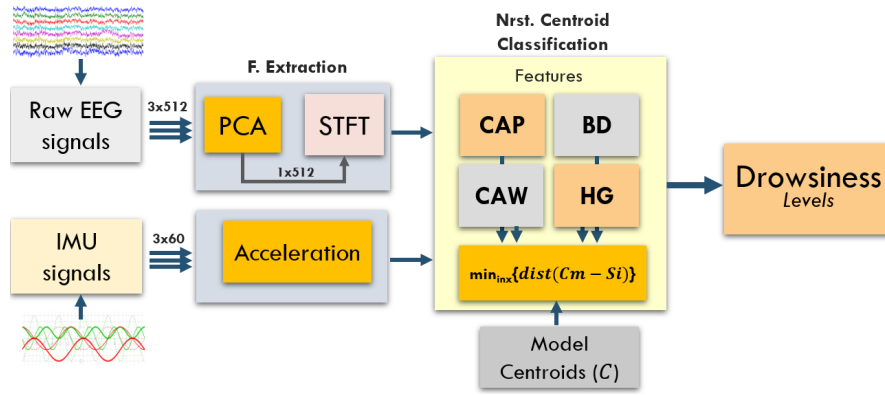
that operates at nominal voltage supply of 1.8V and 2.5V. The difference in energy between the MCUs and the single-core PULP platform at the nominal supply voltage is mainly given by technology gap, different implementation strategy and architectural complexity, which leads to 8.6x to 45.2x lower energy consumption.

More interesting is the exploitation of parallel near threshold computing on the PULP platform (PulpV3), leading to a further improvement of 3.4x in performance with respect to sequential processing, and to an improvement of 12.1x and 63.3x in terms of energy consumption with respect to commercial MCUs. At a system level, this approach has the potential to improve the energy efficiency (to our initial deployment on ARM) by 7x, extending the battery life to 46 hours.

**Table 4.4:** Comparison between different platforms for real-time operation.

MCU	A. Apollo	STM32F407	1C PULP	4C PULP
No. of Cores	1	1	1	4
RT Freq [MHz]	31.68	31.68	45.59	11.76
Vdd (V)	1.80	2.50	0.48	0.45
Pw Dens [W/MHz]	115	600	10.27	27.64
Power [mW]	3.64	18.99	0.42	0.30
Energy[J]	7.28	37.97	0.84	0.59

The real-time frequency (RT freq) corresponds to the frequency required to achieve real-time operation. Ambiq Apollo features a lower power consumption but it does not achieve the required frequencies (i.e. max frequency is 24 MHz). The 4-cores PULP MCU delivers best energy efficiency while meeting the real-time constraint of the application.



**Figure 4.8:** Block diagram of the feature extraction and classification for the drowsiness detection with BioWolf.

#### 4.1.8 Drowsiness Detection With BioWolf

The results obtained when porting the kernel functions of the previously introduced work motivated an actual implementation on BioWolf. As before, a 5-level drowsiness detection system was designed based on a similar feature extraction while employing the Nearest Centroid Classifier to output the alarm levels.

##### 4.1.8.1 Feature Extraction

EEG data are acquired using a three-channel configuration. The first step of the pre-processing is to merge all the information contained in the channels using the Principal Component Analysis. PCA is a linear transformation which represents data into a new reference system by maximizing the variance contained in the original signals [99]. In this case, we reduce the dimensionality of the input matrix to a single vector, retaining up to 90% of the information. From this 1-D vector, we extract four features used for the drowsiness detection.

The first two features are based on the power of the signal in the alpha band (i.e. 8-13 Hz). The first, called Cumulative Alpha Power (CAP) accounts for bursts of alpha waves, which are increasingly present as drowsiness raises [100]. CAP is calculated by computing the alpha-band power spectrum over a moving-average filter with a large size ( $n = 2k$ , or 2 seconds). The second feature, called Constant Alpha Waves (CAW), aims to detect the closure of the eyes that is characterized by a constant presence of alpha waves with a relative large amplitude ( $\tilde{100}\mu\text{V}$ ) and it is extracted by evaluating the signal power on the alpha band over a shorter moving-average window ( $n=500$  or 500 ms).



The third feature, namely Blink Duration (BD), is also extracted from the EEG signal trace, but it derives from an artifact related to extra-ocular muscle movements. The final feature, called Head Gesture (HD) is obtained by processing data from the IMU sensor. These two features are extracted using the techniques introduced in Section 4.1.2. Thus, a further description is omitted.

#### 4.1.8.2 Training and Classification

The classification is based on the Nearest Centroid Classifier, where the output depends on the distance from the average center of the labeled data (centroids). The learning process relies on finding the positions of such centroids, and it is performed automatically using the k-Means algorithm. K-Means clusters data by an iterative process where it first assigns a class label to all observations, as a function of the distance from the current class centroids and then reassigns such centroids to a new position in function of the average position of the recently labeled data. We run this on experimental data from three subjects simulating four different drowsy conditions ( $k = 4 + 1$ ) to create the model of the system. A summary of all the features and drowsiness conditions is presented in Table 4.5.

The initial positions of the centroids play a crucial role in the clustering. We compared two methods, the first one based on random assignment, the second one providing a portion of labeled data for fixing the initial centroid positions. This last method has been demonstrated to help increase the robustness of this approach [102], which is confirmed by our results presented in the next section. Fig. 4.8 summarizes the complete process of feature extraction and classification.

**Table 4.5:** System Drowsiness Levels

Levels of alarms with meeting criteria		
Level	Description	Parameters
-	Fully awake subject (morning test)	$BD < 0.5s$ & $AW < Th1^a$
1	Increase of the blink duration	$BD > 0.5s$
2	Increase of the alpha wave activity	$AW > Th1^a$
3	Head gesture	IMU
4	Closure of eyes	$AW > Th2^b$

$Th1^a$  AW over given threshold over a time window of 20 seconds.

$Th2^b$  AW over given threshold over a time window of 3 seconds.

Data samples were stored after reaching some drowsiness criteria (previously acknowledged [101]) for each level. Physiological parameters were evoked by exposing the subject to drowsiness conditions or simulation.

### 4.1.8.3 Experimental Results

As introduced before, the NCC classification depends on the distance of the observation to the class centroids, which contains information about the model of the classifier. K-means is used to determine the centroids' positions from the data collected from three test subjects simulating four different drowsiness levels. Nevertheless, the model can perform adversely if the initial conditions are not chosen adequately during the training phase. Using two-fold cross-validation, we test two different training strategies. The first trains the model starting from randomly-initialized centroids, while in the second, we provided a certain amount of labeled data to help the clustering. Table 4.6 demonstrates that the second approach leads to significant advantages. This is further confirmed by the online testing results reported in Table 4.7, showing an average accuracy of 83%, with L3 and L4 above 90%, ensuring the detection of the most critical drowsiness states.

During real-time classification, each 8ms a new window of data is elaborated (8 samples overlap). The cluster elaborates the entire processing chain in less than 1ms working with an operative frequency of 100MHz at 0.8V. As a consequence, the total power of the device is around 6.17 mW, and its the contribution of the three main chips mounted on the board. The AFE is responsible for 36% of the power, whether the Nordic MCU employs 43%. The remaining power consumption derives from Mr. Wolf, and it is the result of the parallelization, the optimizations, and power management techniques. Significant energy savings were achieved through the FFT as it took full advantage of the parallel architecture achieving the ideal speedup, only limited by the FPU of Mr. Wolf [103]. During the processing, only the required cores of the cluster are clocked up, avoiding energy loss. When the MCU is in idle, we power off the cluster and part of the SoC (sleep mode) to further reduce the power consumption. As a result, the system

**Table 4.6:** K-means Offline test per level

Method	L0	L1	L2	L3	L4	Avg
<b>Random</b>	0.54	0.65	0.40	0.85	0.95	0.67
<b>Pre-labeled</b>	0.73	0.71	0.71	1.00	1.00	0.83

**Table 4.7:** Online accuracy per level for the NCC

Method	L0	L1	L2	L3	L4	Avg
<b>T1</b>	0.83	0.72	0.74	1.00	0.95	0.84
<b>T2</b>	0.79	0.75	0.64	0.95	0.93	0.81
<b>T3</b>	0.82	0.74	0.70	0.96	0.94	0.83
<b>Avg</b>	0.81	0.74	0.69	0.97	0.94	0.83

delivers up to 60h of autonomy with a 200mAh battery, which can be further extended using the energy harvester subsystem.

#### 4.1.9 Discussion

Drowsiness detection poses significant challenges since the clinical parameters to quantify the drowsiness are not clearly defined, and the systems to detect drowsiness and fatigue level should target unobtrusiveness and energy efficiency.

This section introduced a drowsiness detection system based on sensor fusion techniques implemented initially on a commercial low-power embedded processor. Many of the current implementations consider only individual features, while fatigue and drowsiness symptoms are highly variable among the population and are exhibited through different symptoms.

The embedded systems presented previously tackle these issues by providing a multi-feature approach to detect five levels of drowsiness from EEG signals and an IMU sensor with a detection accuracy in line with the SoA.

The online and real-time implementation of the system implemented on the ARM Cortex-M4 reached 6 hours of life with a 200mAh Li-Ion battery. This initial work also highlighted the benefits of implementing the kernel algorithms on a PULP architecture (PulpV3), showing that an energy boost of 63x can be achieved that can extend the battery life of the system up to 7x to out original implementation.

Further efforts have been made to pursue an actual implementation, which was possible thanks to the development of BioWolf. Apart from the improvements in terms of intrusiveness, the new system, compliant with the SoA (as the same as the former), showed significant advantages in terms of power consumption, reducing up to 17x the power envelope required (6mW). Furthermore, the low CPU utilization of the system suggests that more advanced ML can be integrated to improve the detection accuracy, which is the scope of future work.

## 4.2 SSVEP-Based BCI

A growing trend in Human-Computer Interaction (HCI) is to integrate computational capabilities into wearable devices, to enable sophisticated and natural interaction modalities. Acting directly by decoding neural activity is a very natural way of interaction and one of the fundamental paradigms of Brain-Computer Interfaces (BCIs) as

well. Following this trend, this section presents an online and real-time application of BioWolf for BCI spelling based on Visual Evoked Potentials through Canonical Correlation Analysis (CCA). The system offers performance comparable with those achieved by state-of-the-art non-embedded systems with an ultra-low-power budget and an unobtrusive sensor interface. The following introduces details about the feature extraction and classification mechanism as the same as details of the optimizations to keep the power consumption below 10mW.

#### 4.2.1 Introduction


First BCI spellers were based on P300 (ERP), an EEG response elicited by presenting a relevant stimulus among non-relevant ones. For example, researchers in [104] presented a BCI capable of detecting 36 different target stimuli associated with the letters of the alphabet and some symbols, leading to the overall performance of 0.17 bits per second. Improvements introduced in [105, 106] led to 0.45 and 0.40 b/s. Nevertheless, these results are not compliant with nowadays standards, resulting in a user that refuses to adopt such a system, especially if not impaired.

Steady-State Visual Evoked Potential (SSVEP) is another BCI paradigm that has been used in more recent works with considerable success [27–30]. Processing requires identifying the frequency (and possibly the phase) of the SSVEP signal to determine which stimuli evoked it.

The SSVEP paradigm is attractive due to its higher signal-to-noise ratio (SNR) in comparison with ERPs, being significantly more immune to eye-related and electrode shifting artifacts when a proper frequency band is used [107]. If relying only on frequency information, SSVEPs present two significant advantages with respect to ERPs, namely, no synchronization, and no training required.

Feature extraction for SSVEP can be performed using simple techniques. An early example is found in [108], where the authors designed and implemented a BCI to help users to input phone numbers based almost entirely on FFT-based Power Spectral Density (PSD) analysis. Nevertheless, the approach based on canonical correlation analysis has demonstrated to deliver better performance [31]. Some attempts [109] have been done to use the CCA for real-time classification, achieving, nevertheless, an ITR lower than 1 bps, with a bulky setup.

The lesson learned is that the development of a high-performance wearable platform for BCI spelling is still an open challenge. Although some systems target a portable



**Figure 4.9:** Average CCA correlation of SSVEP responses for different stimuli (x-axes) calculated with different reference signals (y-axes).

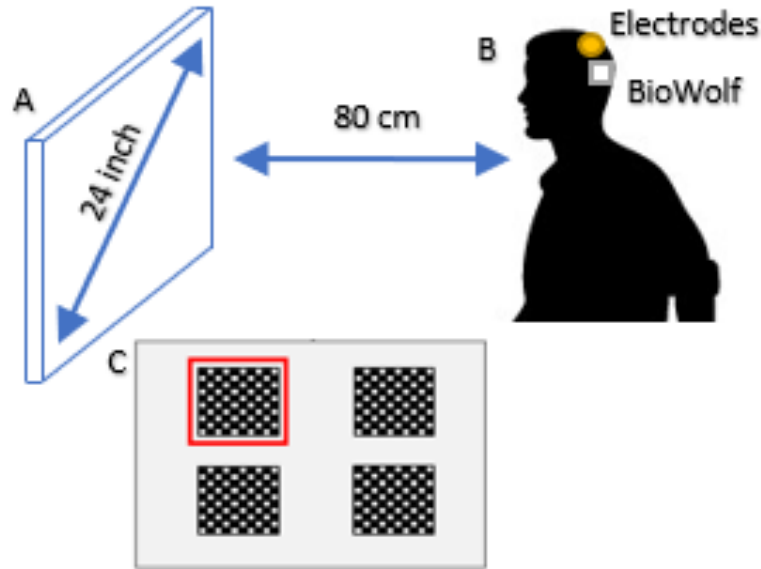
setup, (e.g., a tablet [60] or a smartphone [62]), they achieve poor performance (ITR) while running in intrusive and power-hungry platforms.

To provide improvements in these aspects, we designed a wearable system for brain-computer interaction, relying on a minimally intrusive setup composed of three zero preparation electrodes, which achieves an ITR in line with SoA ( $>1\text{bps}$ ) while running entirely on a wearable node.

#### 4.2.2 BCI System

The proposed BCI includes two main functional blocks, namely, the SSVEP stimulation screen (stimuli presentation) and BioWolf (acquisition and processing). The stimuli are presented on a 24-inch LED (60 fps) display, placed approximately 80 cm from the subject. The stimulation layout is composed of four black and white  $10 \times 10$  square checkerboards [110] arranged in a  $2 \times 2$  pattern at equidistant positions, each displaying a different frequency-coded stimulus, where a single checkerboard occupies 20% of the screen.

Stimuli modulation on each checkerboard is performed using the sampled sinusoidal stimulation method [111], with the contrast of each checkerboard adjusted as follows:



**Figure 4.10:** BCI setup during experimentation.

A - Stimulation screen at 80 cm of the subject. B - Approximate location of the BCI elements. C - Stimulation as seen by the subject.

$$\text{Contrast}(f, \phi, i) = A \cdot \sin(2\pi fi/F_r) + A$$

where  $i$  indicates the frame index,  $A$  the initial amplitude,  $f$  the frequency of the stimulation, and  $F_r$  is the refresh rate of the screen. To help visual fixation, we included a diagonal cross on each checkerboard. All the textures are generated using Psychtoolbox 3.0.10 for Windows in Matlab R2018b.

The frequency targets used are in the range of [6 7.5], with a  $\Delta F = 0.5$  Hz and have been chosen in previous experimentation, including a larger set of frequencies as introduced in Fig. 4.9, where it is evident that the lower half spectrum offers higher performance. To minimize any interference from the theta band and other artifacts [112], we tested several CCA evaluation window sizes, where a window of 2 seconds offered the best ITR values.

During experiments, the subjects are asked to fixate on one of the four checkerboards, as indicated by a red cue. During the stimulation, all four checkerboards are displayed for five seconds. This process is repeated (cue and stimulation) to present twice all frequency targets, for a total of eight trials. For all experiments, BioWolf samples EEG signals at 500 SPS, using only three electrodes, located at POz, PO5, and PO7. The number and location of these electrodes have been fixed during previous empirical experimentation and provide similar performance from a full configuration (eight electrodes at the visual cortex region). The complete BCI setup is summarized in Fig. 4.10.

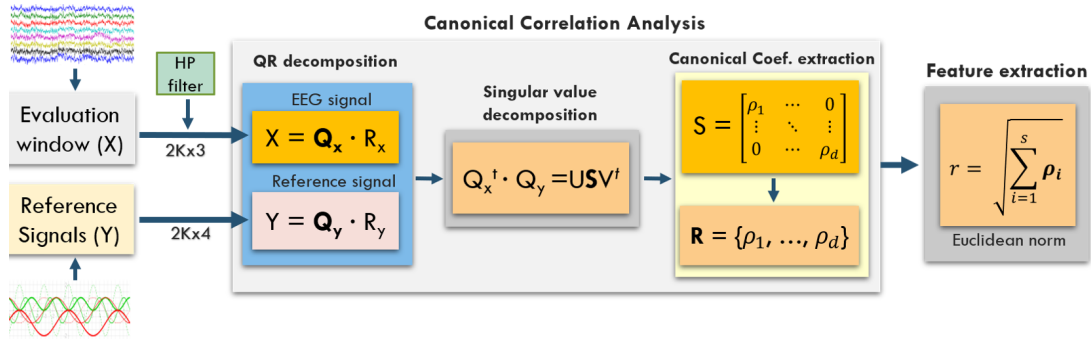


Figure 4.11: Block diagram of CCA algorithm and implementation.

### 4.2.3 Signal Processing

As mentioned in Section 4.2.1, the approach based on canonical correlation analysis is widely accepted as the SoA paradigm [31] for SSVEPs. CCA calculates the *canonical coefficients*, i.e., the maximal correlation between two sets of multidimensional variables. In this work, we used the Golub-Reinsch algorithm [113] to implement the CCA, due to its computational efficiency and scalability, particularly useful for resource-constrained platforms.

Fig. 4.11 shows a block diagram of the signal processing chain. A new window is first band-pass filtered with a 5-taps low pass (LPF) and 2-tap IIR HPF (BW=4-16Hz) and then downsampled. Filtering allows retaining only relevant information and reduces the overall signal noise while downsampling reduces the number of samples that need to be processed by the CCA. The optimal downsampling rate (10 $\times$ ) is derived from experimental data, as in a previous work [114].

The filtered (and downsampled) EEG signals are then used as inputs for the CCA block.  $X_{N_s \times n}$  is EEG data matrix ( $N_s$  samples  $\times$   $n$  channels), while the columns of  $Y$  contain the sine and cosine of the fundamental and second harmonic of the target frequency. CCA is calculated on an  $N_s$  sample window on the  $n$  channels.

The computational steps are the following:  $X$  and  $Y$  are QR-factorized via Householder reduction to extract the orthogonal matrices  $Q_x$  and  $Q_y$ . Finally, the Single Value Decomposition (SVD) is applied to the matrix  $A = Q_x^T Q_y$ . The first part of the SVD reduces the input matrix in a bi-diagonal form via Householder rotations while the second part executes the matrix diagonalization of the bi-diagonal form via symmetric QR decomposition (Givens Rotation). The diagonal matrix  $S$  of the SVD decomposition holds the set of canonical coefficients  $\rho_1 \dots \rho_d$ .

Target classification is finally performed using a threshold over the Euclidean norm between the canonical coefficients. If the correlation of a given frequency exceeds the

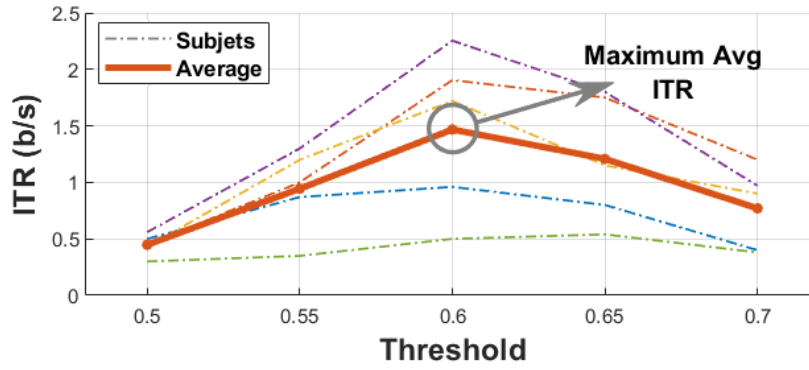


Figure 4.12: Average ITR results calculated with different thresholds values.

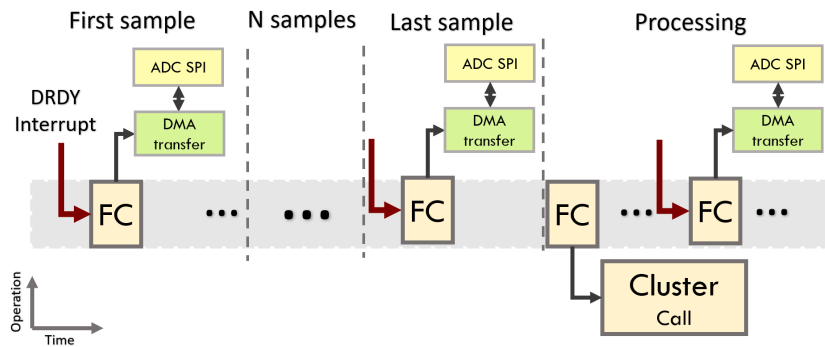


Figure 4.13: Diagram of the signal sampling and execution of the application on Mr. Wolf.

threshold, then the BCI output is the corresponding frequency-coded class. This parameter is chosen to maximize the average ITR over all the subjects, which is calculated for the asynchronous BCIs as in [115]:

$$ITR = \frac{1-P_r}{d_{avg}} \left( \log_2 N_f + (1 - P_w) \log_2(1 - P_w) + P_w \log_2 \left( \frac{P_w}{N_f - 1} \right) \right)$$

where  $P_r$  is the probability of non-detected stimuli,  $P_w$  is the probability of incorrect detected cases,  $N_f$  is the number of target stimuli, and  $d_{avg}$  is the average delay or latency of the system in seconds.

Fig. 4.12 presents different ITR values calculated for several thresholds, where it is noticeable that the value of 0.6 maximizes the average ITR (solid red), and thus, it was selected for the final implementation of the BCI. The same trend can be found for each subject (dashed lines).

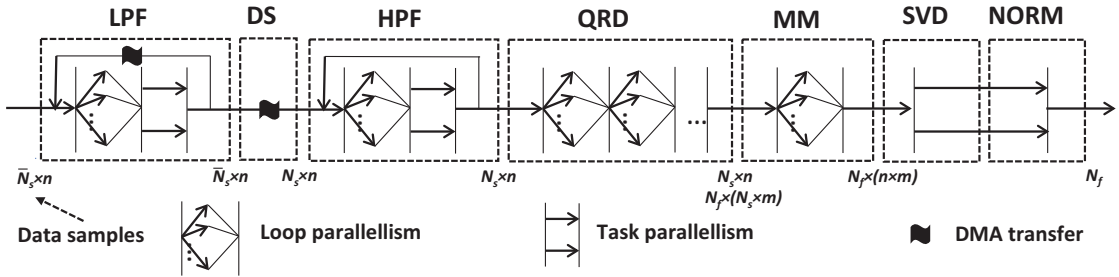


#### 4.2.4 Parallel Processing and Optimizations

This section describes the optimized implementation of the algorithm on Mr. Wolf SoC. Starting from an analysis of the computational requirements of the application, we describe how the blocks of the algorithm presented in Section 4.2.3 have been efficiently implemented on the parallel cluster of Mr. Wolf SoC. The application was implemented using the 32-bit floating-point format, to deal with the high dynamic range of the algorithm, and exploiting the shared floating-point units available on Mr. Wolf SoC. A diagram of the signal sampling and execution of the application on Mr. Wolf is presented in Fig. 4.13, where each block represents the active time of the internal blocks of Mr. Wolf (not in scale). The complete MCU is usually in deep sleep mode. Only when it is required, specific modules are activated. For instance, the sampling of data employs the FC to configure a DMA transfer to read data from the SPI. Once set, the DMA transfer starts, and FC is put back into sleep mode. When the evaluation window is full, the FC calls the cluster to execute the classification and goes into sleep mode. During the cluster processing, the FC is again triggered to fill the next evaluation window in a new buffer.

A breakdown analysis at execution time highlights that the initial LPF covers over 37% of the total cycles, the downsampling filter less than 1%, the HPF about 3%, and finally, the CCA covers the remaining 59%. As discussed in the previous subsection, CCA can be further decomposed into QR decomposition (47%), matrix multiplication (10%), and SVD followed by Euclidean norm (2%).

To parallelize efficiently the application we exploited two different techniques: *loop-level parallelism*, which consists of splitting an iterative workload (i.e. *for loop*) equally distributing the loop iterations among the available cores; *task-level parallelism*, which assigns a region of the workload (a *task*) to a specific core whenever a condition on its numeric identifier is met. In other words, loop-level parallelism splits the data space of the problem into sub-regions, and assign each space to a processor, following a Single Instruction Multiple Data (SIMD) model, where each processor executes the same instructions on different data. On the other hand, task-level parallelism assigns an entirely different task to each core unrelated to one of the other cores, following a Multiple Instruction Multiple Data (MIMD) approach. The cluster architecture of Mr. Wolf allows executing efficiently both computational models, thanks to the tightly coupled data sharing mechanism based on the word-level interleaved multi-bank L1 memory and the availability of hardware-assisted synchronization. In general, a solution based on loop-level parallelism provides higher speed-ups compared to task-level parallelism since it is easier for the programmer to enforce workload balancing among the available cores. However, the structure of some algorithms does not allow the adoption of



**Figure 4.14:** Parallelization schema of the processing pipeline highlighting the execution flow and its branches. This diagram reports the size of data samples and the parallelization approach (loop-level or task-level) for each computing block.

loop-level parallelism due to data dependencies among the iterations: In these cases, task-level parallelism is the only viable solution. Fig. 4.14 shows the parallelization scheme adopted in each component of the processing toolchain.

The LPF performs its computation on data structures which are allocated in the L2 memory since the total amount of data does not fit the L1 memory level. To hide the higher access latency of the L2 memory, we have designed a *double-buffering* technique using the cluster DMA and L1 buffers for intermediate results. More in detail, processing of input data is divided into  $c$  chunks of size  $N_{chunk} \times n$  with  $N_{chunk} = \bar{N}_S/c$ . While a set of input/output buffers is accessed by processing cores for computation, a distinct set of buffers is used as destination (i.e., the next chunk to compute) or source (i.e., the result of the previous computation) of a DMA transfer. Since the backward computation is affected by data dependencies, we applied task-level parallelism only to this part, assigning a core to each channel (from 0 to  $n - 1$ ).

The downsampling filter (DS) reads input data from the L2 memory and writes its result in L1. This filtering stage is performed using the 2D capabilities of the DMA, which is initially programmed by core 0 and then does not require any additional contribution by the processing cores.

The HPF adopts the same parallelization strategy of the LPF. However, its impact on the total execution time is almost negligible (3%). In this case, both input and output data fit the L1 memory area, and double-buffering is not required. The QR decomposition of the target frequencies is performed offline, so this block is not included in the figure.

The final part of the processing chain includes SVD followed by Euclidean normalization (NORM). Since data dimensions are lower than the number of cores, the computation of SVD does not expose enough work-load for multiple cores; moreover, data dependencies among the iterations do not allow to apply loop-level parallelism on

outer loops, while the work-load of inner loops is negligible. Considering these limitations, we applied task-level parallelism to the code block SVD+NORM assigning a core to each frequency (from 0 to  $N_f$ ).

The parallel speed-ups achievable from SVD and backward computations in LPF and HPF are limited by construction if the number of cores is higher than the number of tasks ( $N_f$  and  $n$ , respectively). As highlighted in the previous paragraphs, this design choice is forced by the data dependencies inherent in the algorithms. Moreover, QRD, LPF, and HPF are heavily FPU-intensive, and their parallel speed-up is further limited by the number of shared FPU available in the architecture.

#### 4.2.5 Experimental Results

Five healthy subjects (aged 25-40 years) with normal or corrected-to-normal vision participated in the offline and online experiments. All participants reported no history of neurological or psychiatric disorders and provided written consent to participate in the experiment. Offline experiments have been performed before the online tests to optimize critical parameters of the systems and maximize the ITR (electrodes count and location, frequency intervals, and evaluation window size) as mentioned before. The same group also participated in the final online experiments that were used to validate the system performance (accuracy, average latency, and power consumption). Our subjects have previously participated in SSVEP-based tests, and hence, training was omitted since they were already familiar with the experiments.

The total accuracy is calculated as the ratio between the number of the correctly classified points over the total, i.e., each classification output of the system (every 100ms) is compared with the expected label, including rest and stimulation states (hence, on

**Table 4.8:** Online results for 4 stimuli.

Subject	Tot accuracy	Trial accuracy	Latency [s]	ITR [b/s]
S1	0.81	0.87	1.42	0.96
S2	0.81	1	1.05	1.90
S3	0.81	1	1.16	1.72
S4	0.87	1	0.88	2.25
S5	0.82	0.87	1.80	0.5
<b>Average</b>	<b>0.82</b>	<b>0.95</b>	<b>1.26</b>	<b>1.46</b>

The total accuracy refers to the classification performance point-to-point with respect to the expected label (a new classification point is available every 100 ms). The trial is instead asserted as correct if the classification output is equal to the expected when there is stimulation. All the values of ITR are calculated assuming an asynchronous BCI.

a point-to-point basis). The trial accuracy is calculated instead as the ratio between the number of correctly classified trials over the total number of trials (a test contains eight trials). A trial is considered correct if the first classification, after the onset of stimulation, matches the expected one.

The trial accuracy is indeed more applicable for real scenarios, but we also decided to provide the results of the total accuracy since this gives insights about how good the system is to detect the rest states.

Table 4.8 summarizes all the previously introduced metrics for our online system. The average total and trial accuracy of the system are 0.82 and 0.95, with an average output latency of 1.26 seconds. As a result, the average ITR is above 1.4 b/s, proving that our wearable and embedded BCI achieves comparable performance with respect to SoA systems [29, 116, 117] while running entirely on a highly wearable low-power platform.

To verify the long-term performance of the system, we performed an additional test, including only one test subject that in addition to the SSVEP test (as in Section 4.2.2), also performed an alpha wave test (eyes closed/open). In an initial step, we registered a baseline minutes after setting up the BCI. We have repeated these experiments after one hour of usage. The CCA of the signals for both tests is presented in Fig. 4.15-A, where, at first glance, there is no appreciable difference (baseline as a solid line and test after one hour as a dashed line). A closer analysis shows that there is a small deterioration in ITR ( $< 5\%$ ) that can be considered negligible and can be associated with physiological parameters of the user rather than to the BCI itself. In Fig. 4.15-B we also present the results for the alpha wave test that is in line with the initial findings.

#### 4.2.6 Power Consumption

**Table 4.9:** Energy/Classification per Kernel Function

Kernel	Mr. Wolf Seq		Mr. Wolf 8-cores	
	cyc(k)	E( $\mu$ J)	cyc(k)	E( $\mu$ J)
LP-Filter	130.29	25.92	45.34	9.02
Downsampling	1.00	0.19	1.00	0.19
HP-Filter	111.85	2.22	3.46	0.68
CCA	210.63	41.91	47.80	9.40
<b>TOTAL</b>	<b>453.77</b>	<b>70.24</b>	<b>97.6</b>	<b>19.29</b>

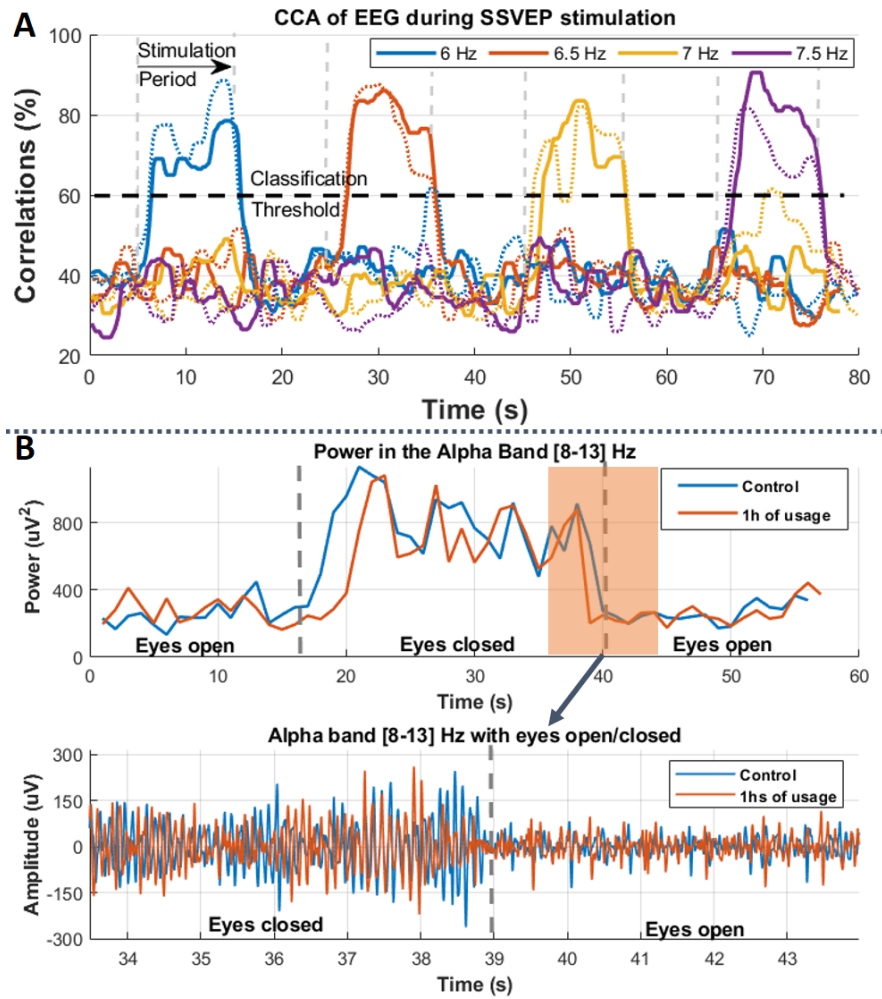
The energy/classification is calculated considering and op. frequency of 100MHz@0.8v.

Table 4.9 presents more detailed information for the energy consumption when running the embedded application. The heaviest processing in our BCI implementation comes from the CCA and the low-pass filter, requiring up to 47k and 45k cycles per output when employing all the cores of the cluster. The downsampling and high-pass filtering offer negligible contributions to the final energy envelope. The benefits of the parallelization are more evident when examining the ratio between the total energy of the sequential and parallel processing of the kernel functions at the bottom of the table, which denotes a  $3.64\times$  energy saving. Further savings are achieved by toggling the power state of Mr. Wolf's fabric controller and cluster when possible. As a result, the total power envelope of Mr. Wolf is around 1 mW. The energy required by the other digital blocks (nRF52840 and digital interface of the AFE) is dominated by the energy required to transmit the result of the computation and is about 1.26 mW (mostly to keep active the BLE link), while the power consumption of the ADC is about 4 mW, that includes the sampling of three EEG channels. As a consequence, the final power envelope is about 6.31 mW, guaranteeing up to 38 hs of battery life. From these measurements, we can observe two remarkable results. The first one is that the system is capable of processing data at a power consumption significantly lower than that required for the bare streaming of data (3.16 less power), with obvious advantages in terms of efficiency of both the wearable node and of the full system. Secondly, the energy efficiency of Mr. Wolf and of the custom implementation of the algorithm make the power budget of the system largely dominated by the acquisition and Analog-to-Digital conversion of data (64%).

We also studied the battery life improvements provided by the EH subsystem present in BioWolf (introduced in Section 3.2.5). Under ideal conditions, the EH subsystem is capable of extending the battery duration of a factor of two, up to 76 hours. In principle, if the application allows for a reduced setup (e.g., using only 1 channel, with a power consumption of 3.76 mW), self-sustainability can be achieved in outdoor environments. However, it should be pointed out that a bright environment might also reduce the SSVEP effect and the performance of the system.

#### 4.2.7 Discussion

This work presented a BCI system based on the canonical correlation analysis for the classification of SSVEPs while running on BioWolf. The system provides SoA performance with a power envelope suitable for wearable applications. The device guarantees up to 38 hs of battery life, which can be translated into 14 days of operation when using the device 2 hs/day. Table 4.11 present a comparison with the SoA for embedded SSVEP classification. The implementation with Biowolf offers an ITR higher than most



**Figure 4.15:** BCI testing after one hour of usage. A - CCA test, where the solid lines denote the CCA output of the SSVEP response on a test performed minutes after setting up the BCI system. The dashed lines denote the CCA output of the SSVEP response after one hour of device wear. B - Alpha waves test. The topmost figure of this section presents the alpha band power of a 55-second trial, where the blue line shows the PSD response from EEG signals acquired on a control test minutes after setting up the device, while the red line corresponds to a test after one hour of device wear. The bottom figure denotes the signal in the time-domain of the highlighted (in red) region of the first figure.

of the listed work, being only 0.10 b/s off from a previous work[114] implemented with the ARM-based BCI described in Section 3.1. Nevertheless, the efficiency offered is more than 15x higher (for the processing unit) with a power budget of 4.35x lower.

To complement the advantages in terms of energy efficiency, we have also demonstrated how the energy harvesting subsystem can provide self-sustained operations when working in reduced mode. The compactness and ease of use of our platform allow it to be seen as an accurate, self-contained wearable device for day-to-day BCI applications. Still, since the implemented algorithm uses less than 5% of Mr. Wolf computational resources, a more demanding complex BCI algorithm can be accommodated (artifact

removal techniques based on blind source separation, advanced machine learning, or Convolutional Neural Networks) as the same as a more significant number of channels.

### 4.3 ERP-Based BCI

Sensory evaluation is used to assess the consumer acceptance of foods or other consumer products as to improve industrial processes and marketing strategies. The procedures currently involved are time-consuming, requiring a statistical approach from measurements and feedback reports from a large set of evaluators under a well-established measurement setup.

In this section, an automatic system for sensory evaluation through Event-related potentials (ERPs) based on BioWolf is presented. Being the outcome of the processing of visual stimuli that can be modulated by the emotional state of the subject, ERPs can be employed to assess the perceived quality of food. This implementation allows narrowing the number of evaluators since errors related to psychological factors are by-passed. The preliminary results on measuring ERPs presented demonstrates the feasibility of this approach towards a fully embedded system.

#### 4.3.1 Introduction

Sensory evaluation is defined as the scientific method used to give a quantitative measure to the appearance or flavor of a food product as perceived through the senses of sight, taste, smell, touch, and hearing [118].

**Table 4.11:** SoA embedded implementations for SSVEP classification.

	Footprint [mm x mm]	Number of Channels	Connectivity	System Power [mW]	SoC	Performance [MFlop/s] <sup>1</sup>	Efficiency [MFlop/s/mW]	ITR
Lin [60] :	n.a.	2	BT	5000 <sup>2</sup>	Qualcomm Snap- dragon 800	9200	1.8 <sup>2</sup>	0.56
Chi [62] :	n.a.	8	BT	1000	TI OMAP 2420	434	0.5 <sup>2</sup>	0.44
Chai [57] :	36x36 <sup>3</sup>	2	2.4GHz RF	>3300	Nordic nRF24	16 <sup>4</sup>	1 <sup>4</sup>	0.45
Salvaro [41] :	85x50	8	BT2.1	27.50	ARM- STM32F407	168	1.2	1.57
<b>This work :</b>	<b>40x20</b>	<b>3<sup>5</sup></b>	<b>BLE</b>	<b>6.31</b>	<b>Mr. Wolf</b>	<b>1400</b> <sup>6</sup>	<b>18<sup>7</sup></b>	<b>1.47</b>

<sup>1</sup> Ideal (i.e. considering maximum operating frequency x # of FPU's).

<sup>2</sup> Estimated.

<sup>3</sup> The complete system is composed of multiple embedded devices.

<sup>4</sup> Fixed-point.

<sup>5</sup> 8-Channel Capable.

<sup>6</sup> Considering two floating-point operations for fused multiply and accumulate. Cluster 1.1V, 350 MHz.

<sup>7</sup> Operating point: 0.8V, 100 MHz.

In the food industry, sensorial analysis is actively used to understand the target market and to optimize the effort and investment during product development. It has become an irreplaceable tool to determine the success of a product as consumers are mostly driven by personal sensorial feedback and by the brand, rather than other essential features like nutrition elements or convenience.

A sensorial analysis is performed through affective and analytical techniques. The former requires large consumer panels to answer a questioner after tasting. Analytical methods are specific tests done by trained experts and can be used only to determine if products are different, or if a food variety highlights a selected characteristic more than another [119].

Psychological and sensorial feedback from stimuli produced by food products play a fundamental role also in the ordinary industrial transformation of food (e.g. fruit and vegetables) as a quality assessment during the whole transformation chain, usually done by a manual screening of the operators.

Therefore, sensory evaluation is an integral part of the industrial food process and the success of a product. Nevertheless, two significant factors limit this method: time and human error/bias.

On the other hand, ERPs are one of the most common techniques to study emotion processing by visual stimuli [120]. Specifically, Late Positive Potentials (LPPs), a brain response visible after 400ms from the stimuli presentation, can be modulated by the emotional intensity of the stimulus [121], where stimuli with a stronger emotional response (to a neutral stimulus) elicit a larger LPP.

In this scope, we propose to collect the signal of the perceived quality of the food from Event-related potentials (ERPs-LPPs) directly and to process them to have an objective evaluation of the sensory perception, without any bias introduced by the operator through higher brain functions in an automatic manner. In the following, details about the processing chain are presented as the same as the wearable deployment of the system.

### 4.3.2 Signal Processing

Data processing is performed on Mr. Wolf. Since computation is relatively simple, all processing can be performed on the Fabric Controller (FC), which processes data in double-precision fixed-point representation. It also takes care of exchanging data with the AFE for signal acquisition and with the nRF32832 SoC for communication with the host through BLE. The latter is a bidirectional communication as the host needs to



provide synchronization signals (*triggers*) to inform the device of the instant at which the visual stimulus is presented to the subject. When the acquisition of a full trial is completed, the device sends the reconstructed ERPs to the host.

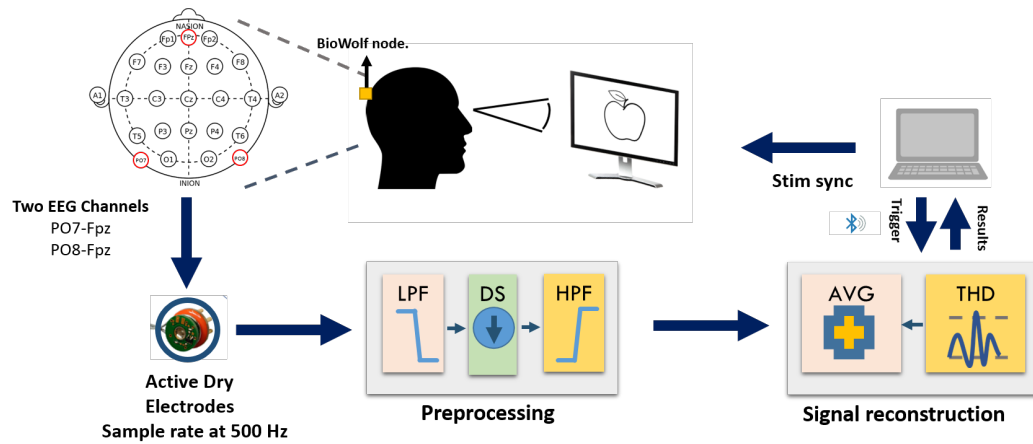
The processing steps for reconstruction of the ERPs are the following:

- *Preprocessing:* the first operation is the computation of the average of data acquired from the two differential channels (PO7-Fpz and PO8-Fpz) at 500 SPS. Data is then filtered with a Finite Impulse Response (FIR) low-pass filter (LPF) with 100 taps and -3dB corner frequency of 30 Hz. This allows to down-sample the data by a factor of 5, reducing the computational load of the filtering steps as both the outputs of the LPF and the subsequent High-Pass Filter (HPF) can be computed at a reduced sampling frequency of 100 SPS. FIR filtering is preferred to IIR despite the increased computational burden because of its linear phase, which allows minimizing distortion in the reconstructed ERP. Since we are interested in late potentials in the ERP, high pass filtering needs to be performed at a very low corner frequency of 0.25 Hz. This requires the use of an extremely high order filter (1000 taps), which dominates the computational burden of the ERP reconstruction algorithm.
- *Epochs reconstruction:* to remove epochs containing artifacts we adopted a simple method based on automatically rejecting epochs containing samples over  $\pm 50\mu V$ . To this purpose, each sample at the output of the HPF is checked at run-time and, if its absolute value is higher than the threshold, the epoch is rejected and every computation for that epoch (including filtering of subsequent samples) is stopped. Information on the time at which the epoch starts is sent from the host via BLE. When the epoch is not rejected, it is averaged with previously accepted epochs to reconstruct the ERP.

Stimuli is presented through images on a 17-inches LCD screen from a PC running Psychtoolbox under MATLAB. Setup and processing steps are summarized in Fig.4.16.


### 4.3.3 Experimental Results

We tested the device on five able-bodied, (aged  $32 \pm 5$  years old) without previous history of neurological disorders. All participants provided written consent to take part in the experiments. The subject was sitting in a dimly lighted room, approximately 80 cm apart from a 17-inches LCD screen. 200 images of good-quality and defective apples (100 per class) were presented randomly on a 17-inches LCD monitor for 1 second,



**Figure 4.16:** Setup and processing steps for the ERP-based BCI.

Images are presented on a 17 inch LCD screen, while EEG is recorded in PO7 and PO8 with reference on Fz. Data is band-pass filtered between 0.25 and 30 Hz and decimated by a factor 5 to 100 sps. Epochs with data above a  $\pm 50 \mu V$  are discarded. The remaining epochs are averaged to provide the final ERPs.




**Figure 4.17:** ERP from subject 1.

Red line corresponds to non-commercial grade apple pictures, blue line to commercial-grade ones. Data is band-pass filtered between 0.25 and 40 Hz, epochs are rejected if EEG amplitude exceeds  $\pm 50 \mu V$ .

separated by a 1-second fixation cross. Event-Related Potentials were measured as the average potential between electrodes PO7 and PO8 versus Fz. Fig. 4.18 presents results obtained through the analysis of ERPs on the five subjects (grand average on top) and on a single subject (bottom).

Averaged responses are flat before the presentation of the stimulus ( $t=0$ ) and present a high correlation until approximately 300 ms (ERP components related to visual processing of the image). After that time, late positive potentials (LPP) components show significant differences in response to commercial-grade vs. non-commercial grade fruit



**Figure 4.18:** Grand average of the ERPs extracted from the five subjects.

Red line corresponds to non-commercial grade apple pictures, blue line to commercial-grade ones. Data is band-pass filtered between 0.25 and 40 Hz, epochs are rejected if EEG amplitude exceeds  $\pm 50$   $\mu\text{V}$ .

images, testifying how responses to the two classes of stimuli are distinguishable, and giving a clear indication on the suitability of this approach for automated quality grading analysis.

To evaluate the performance of the embedded implementation, we set the operating frequency and voltage of Mr. Wolf at 50 MHz and 0.8 V, respectively. Although this frequency can be scaled up to 450 MHz to meet more constrained applications, the processing required in this work is minimal. Thus, a lower operational frequency satisfies the real-time constraints while minimizing the overall power consumption.

The power consumption of the system is the contribution of the active blocks, namely, Mr. Wolf, the ADC, and the Nordic SoC, for a total of 10.88 mW. Data sampling through the analog sections (ADS1298) requires the highest share of the overall power (around 80%). Although the use of this ADC comes at a high power cost, it ensures the required signal quality, also allowing to avoid more rigorous filtering of the signal that could be translated into higher power consumption at the MCU side. The digital section that includes occasional BLE transmission of computation results and synchronization of the trigger, and the data transfers between AFE and Mr. Wolf, represents the 11% of the total power.

Mr. Wolf is responsible for the remaining power, which is the result of power management techniques, such as switching through the MCU power modes. While working in Run Mode, the acquisition and processing of a single sample only require 0.5 ms, and, when in idle, the MCU is put into sleep mode to minimize the power consumption. Since the FC is able to satisfy the real-time requirements of the application, the cluster

remains off all the time. This demonstrates the versatility of Mr. Wolf to work at a low power budget throughout different computational needs.

As a result, our system achieves up to 19 h of autonomy with a 60 mAh battery, which can be further extended up to 20 h and 40 h in indoor (600 lux)/outdoor (10000 lux) scenarios, respectively, using the EH subsystem.

#### 4.3.4 Discussion

This work presented an ERP measurement system featuring BioWolf, a Parallel Ultra Low Power platform, which allows real-time onboard EEG signal acquisition and processing for differentiation of classes of images to rank food quality. The results show a clear separation of two distinct ERPs responses in correspondence to commercial and non-commercial grade apples. The processing of the EEG signal is performed on-line on the Mr. Wolf platform, which provides more than 19h of battery life with a tiny 60 mAh LiPo battery. The preliminary results presented validate our embedded deployment and represent a step forward towards an unbiased automated food quality grading.

### 4.4 Online Learning and Classification of EMG-Based Gestures using Hyperdimensional Computing

The development of wearable sensing technologies and unobtrusive devices is paving the way to the design of compelling applications for the next generation of systems for a smart IoT node for Human-Machine Interaction (HMI). In this vein, the following section presents a hand gesture recognition application, which is a preferred way of interaction in HMI design.

This implementation, based on BioWolf, runs a machine learning algorithm (HDC) in real-time, recognizing up to 11 gestures with a power envelope of 11.84 mW. As a result, the proposed approach delivers up to 35 hours of continuous operation and 1000 hours in standby. The resulting platform minimizes effectively the power required to run the software application, and thus, it allows more power budget for high-quality AFE.

#### 4.4.1 Introduction

The hand gesture is probably the most natural and direct method used by humans to interact with objects, and it has compelling and straightforward applications in many scenarios, as described before.

Decoding human intentions expressed by hand gestures is usually based on two main approaches: visual recognition of hand gestures using computer vision techniques [122], and recognition based on the analysis of the electrical activity of the muscles involved in the gestures [7].

The former is typically able to recognize a large number of gestures [122], requiring, nevertheless, external infrastructure (cameras and power supply) while being sensitive to environmental factors (line of sight and illumination). The latter is based on decoding EMG signal by leveraging techniques including direct control [9], pattern recognition [123] and, deep learning [124].

Such systems require accurate sensory interfaces and high computational capabilities to be implemented on systems with a reduced form factor, due to the intrinsically noisy nature of the EMG signal and on the computationally demanding algorithms required to make sense of the biosignals [125].

Some attempts have been made at a commercial level, such as the MYO [126], an armband that acquires EMG data from 8 differential channels and sends the data collected on EMG to a PC that processes them with pattern recognition techniques, to recognize up to 5 gestures. Such an approach requires a continuous link between the sensor armband and the PC/gateway platform since traditional wearable platforms are not suitable for computationally intensive tasks, such as pattern recognition algorithms.

To move towards fully portable solutions, an approach that is gaining traction is to use an offline bench-top system for the algorithm training and to implement the classification of the EMG signal directly on the wearable node. However, designing wearable integrated systems for the acquisition and processing of EMG signals capable of executing full pattern recognition algorithms in real-time at high energy efficiency is still an open challenge.

Some systems, like the work presented in [127] or [128], rely on high-end ARM CORTEX A8 processors, which can sustain the high computational load but require significant energy, guaranteeing only 0.5 h of operation with a 100 mAh battery.

More efficient solutions, such as [129] and [42] are based on dedicated industrial IoT microcontrollers (i.e., ARM Cortex-M4) and provide up to 10 hours with a 100 mAh LiPo battery.

The lesson learned from this analysis is that the development of HMI wearable devices poses two significant challenges for the digital processing part. First, the power envelope of the digital platforms must be minimized to allow high-quality signal acquisition via an Analog-Front-End (AFE). Second, approaches based on data streaming,

which offloads the signal processing on external platforms, do not scale well because of limited bandwidth and a high energy-per-transmitted bit of wireless interfaces, even with energy-efficient protocols such as Bluetooth Low Energy.

In this work, we take advantage of the features present in BioWolf to implement a novel framework for EMG gesture recognition based on Hyperdimensional Computing [130], a novel pattern recognition framework. We characterize the performance of the system in terms of energy efficiency, showing that, while running the application, the device consumes only 11.84 mW, providing up to 18 hs of operation with a battery life that can be extended through the EH subsystem. The full HMI recognition software runs on the wearable node that employs less than 30% of the total power to acquire and convert the EMG signals. Thus, the remaining power can be employed on power-demanding high-quality AFEs, resulting in an improvement of the overall performance of the system. In the following, details of the implementation are presented.


#### 4.4.2 Signal Processing and Classification

HD Computing algorithm is a brain-inspired approach that computes with points in the HD space (hypervectors) as an alternative to numbers [130]. Hypervectors are considered as (pseudo)random dense binary vectors composed of an equal amount of randomly placed 0s and 1s. Working with binary allows reducing computational complexity by replacing integer computation into binary operations, where Multiplication, Addition, and Permutation (MAP) are substituted by componentwise XOR, componentwise majority, and one-bit circular rotation ( $\rho$ ), respectively.

Features are extracted from the raw signals and mapped into the HD space using Item Memory (IM) and Continuous Item Memory (CIM) [131] matrices. In this work, feature extraction is based on the RMS envelope of the signal. The IM is composed of random orthogonal ( $\perp$ ) hypervectors (i.e.,  $E_1 \perp E_2 \dots \perp E_i$ ) related to the input channels. The CIM contains orthogonal endpoint hypervectors, mapped through discretized values of the input channels. The input values are discretized to be associated with a given hypervector, where  $V_1$  and  $V_K$  are related to the minimum and maximum input values. The intermediate levels are generated by linear interpolation between these two orthogonal endpoints [131].

The HD computing provides two encoders, spatial and temporal. The first captures the spatial information contained in the signal with a componentwise XOR between  $E$  and  $V$  resulting (at instant  $t$ ):

$$S^t = [(E_1 \oplus V_{l(1)}^t) + \dots + (E_i \oplus V_{l(i)}^t)]. \quad (4.1)$$



**Figure 4.19:** Implementation on BioWolf of the HD computing algorithm.

The temporal encoder instead, extracts information about the temporal evolution of the input data through permutation and multiplication of  $n$  consecutive hypervectors generated by the previous encoder.

Thus,  $n$  spatial hypervectors form an  $n$ -gram hypervector ( $T$ ), defined as:

$$T = S^t \oplus \rho S^{t+1} \oplus \rho^2 S^{t+2} \oplus \dots \oplus \rho^{n-1} S^{t+n-1} \quad (4.2)$$

where  $\rho^k$  stands for  $k$  times permutation.


The HD generates different  $n$ -grams for each gesture. These are finally added to create a *prototype* hypervector stored in the associative memory (AM).

During testing, a new sample is encoded into a  $n$ -gram (*query*) hypervector and compared with all the prototype hypervectors in AM through the Hamming distance, where the label associated with the minimum distance corresponds to the classification output. Fig. 4.19 shows a block diagram of the complete system.


As binary hypervectors assume a very high dimension (i.e., 10k-D), for the embedded implementation, they have been compacting them into 32-bit unsigned integer, leading to a conspicuous gain in performance and memory requirements. And, although this representation requires bitwise operations for multiplication, addition, and permutation (MAP) operations (i.e. read/insert bits into a 32-bit word and popcount), thanks to the presence of the RI5CY processor, which offers aggressive performance optimizations, including bit manipulation instructions (builtins), these operations can be performed in a single clock cycle[132], dramatically reducing the computational load on the MCU. Another optimization derives from the exploiting of the parallel programming models through an optimized version of Open Multi-Processing (OpenMP).

#### 4.4.3 Experimental Results

To demonstrate the performance of the system in terms of classification accuracy, we involved ten able-bodied subjects (aged  $32 \pm 5$  years old) without a previous history of neurological or muscular disorders for the experiments. All participants provided



**Figure 4.20:** Gestures used for the testing. Open hand, fist, index, 2-fingers pinch, rest position.



**Figure 4.21:** Average accuracy obtaining by HD computing, using the same data collected by 10 subjects, increasing the number of gestures (from 1 to 11).

written consent to participate in the experiments. The gestures tested in this work are open hand, fist, index, 2-fingers pinch, ok, supination, pronation, number two, number three, number four and rest position as shown in Fig 4.20.

The algorithm is trained for each subject off-line and the AM matrix stored in the L2 memory. The training can also be executed on-chip in real-time, but this is out from the scope of this work. Fig. 4.21 shows the average accuracy results obtained by increasing the number of gestures (from 2 to 11). The accuracy stands between 84.3% and 99.4%, showing that this implementation is suitable for a hand gesture controller [42].

Table 4.12 shows performance in execution time and energy consumption obtained by executing the algorithm on different configurations of the target architecture. The first kernel (RMS) computes the envelope of the raw signals on a circular buffer of dimension



60. It does not require bitwise operations. Hence, the built-ins are not involved. This kernel can be perfectly parallelized on eight cores as each core can extract the envelope from 1 channel. In the MAP+ENCs kernel, the cluster executes the component-wise XOR operation between CIM and IM and the component-wise majority to create the spatial hypervector. This is optimized through the built-ins, obtaining  $2.6\times$  better performance. Moreover, the workload is equally distributed among the cores of the cluster (each core performs the encoding operations on a different portion of the hypervector), showing a gain of  $20.4\times$  ( $7.7\times$  wrt Mr. Wolf 1 core with built-ins). In the last kernel (AM), the *query* hypervector in output from the MAP+ENCs kernel is associated with one of the possible gestures. Here, it is possible to optimize the performance of the component-wise majority and the popcount ( $2.8\times$ ) used for the Hamming distance through the built-ins. The small quantity of work to distribute among multiple cores leads to a saturation of the speed-up. The small gain obtained in this kernel ( $9.5\times$ ) does not impact significantly on the overall performance ( $17.9\times$ ) because of the dominance of the MAP+ENCs kernel.

To evaluate the performance of the architecture, Mr. Wolf’s operating frequency was set to its most efficient operative point, 100 MHz at 0.8 V. Table 4.12 shows results related to the energy consumed for the classification of a new sample. The dominant part of the entire processing derives from the MAP+ENCs kernel with an energy consumption of  $71.9\ \mu\text{J}$ . The optimized version with the built-ins leads to a gain of  $2.6\times$ , which is further improved, exploiting the parallel computing on eight cores ( $13.0\times$ ). The overall energy consumption of the single-core execution is  $81.44\ \mu\text{J}$ , further reduced by the introduction of built-ins ( $2.6\times$ ). Furthermore, splitting the workload among the eight cores leads to a total energy consumption of  $7.2\ \mu\text{J}$  for a single classification.

While running the application, the total power consumption of the system derives from the contribution of the active blocks, namely, Mr. Wolf, the ADC, and the Nordic Soc, for a total of 11.84 mW. The analog sections (mainly the AFE) is responsible for

**Table 4.12:** HD Computing Execution times on the target architectures, with 10,000-D,  $N=1$ . (Cyc, su) stand for (cycles, speed-up). The total energy/class reported, is the result of the addition of the contribution of these functions without considering the energy during idle periods.

	Mr. Wolf 1 core		Mr. Wolf 1 core built-ins		Mr. Wolf 8 cores built-ins			
Kernel	cyc(k) <sup>a</sup>	E( $\mu\text{J}$ ) <sup>c</sup>	cyc(k) <sup>a</sup>	su <sup>b</sup>	E( $\mu\text{J}$ ) <sup>c</sup>	cyc(k) <sup>a</sup>	su <sup>b</sup>	E( $\mu\text{J}$ ) <sup>c</sup>
RMS	6.82	0.86	6.82	1.00	0.86	0.89	7.66	0.17
MAP+ENCs	569.10	71.91	215.35	2.64	27.21	27.94	20.36	5.55
AM	68.59	8.66	24.19	2.83	3.05	7.23	9.48	1.43
TOTAL	644.48	81.44	246.37	2.62	31.13	36.06	17.87	7.17

<sup>a</sup> cycles per sample, <sup>b</sup> speed-up wrt Mr. Wolf 1 core, <sup>c</sup> 100MHz@0.8V

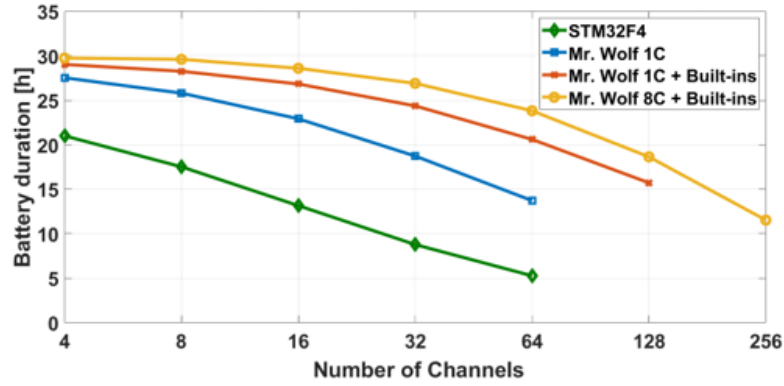
67% of the power consumption, whether the digital section (mostly BLE transmission of computation results, data transfer between AFE and Mr. Wolf) employs 13%. The remaining power consumption derives from Mr. Wolf (SoC and cluster), and it is the result of the parallelization, the optimizations, and several power-management techniques. Data is acquired at a sampling frequency of 1 KHz, and a new data window is elaborated each 8 ms (8 samples overlap). The cluster elaborates the entire processing chain in less than 1ms. During the processing, only the required cores of the cluster are clocked up, avoiding energy loss. When the MCU is in idle, the cluster and part of the SoC remain in sleep mode to minimize the power consumption. As a result, our system delivers up to 18 h of autonomy with a 60 mAh battery, which can be further extended up to 19 h and 35 h in indoor (600 lux)/outdoor (10000 lux) scenarios, respectively, using the energy harvester subsystem. These results are based on the values summarized in Table 4.13, where we also show the current consumption of the system in streaming mode, with up to 9 h of autonomy, and sleep/standby (up to 1000 h). While it is difficult to compare wearable systems directly, it is still noticeable that SoA systems for EMG gesture recognition have a battery life ranging from 3 to 11h [43], [133], [129], independently from the algorithm that is used. As explained above, our architecture is capable of providing around 2x more autonomy with a tiny 60 mAh battery, offering superior performance and unintrusive form factor.

#### 4.4.4 Discussion

In this section, a complete system for hand gesture recognition was presented. The performance of the proposed system, both in terms of execution time and of energy efficiency, allows the design of a smart interface to communicate with objects through the hands by virtue of its highly optimized and versatile architecture, which combines a small solar harvester with an energy efficient and versatile chip, Biowolf can run a pattern recognition algorithm to classify up to 11 hand gestures while ensuring up to 18 h of continuous operation that can be further extended up to 35 h with outdoor illumination.

**Table 4.13:** Current Consumption of the Board Components in the Different Operational States

	DSP Mr. Wolf	Digital Section 1.8v	Analog Section 2.7v	Battery Drain
<b>Sleep</b>	55 $\mu$ A	10 $\mu$ A	10 $\mu$ A	50 $\mu$ A
<b>Streaming</b>	55 $\mu$ A	7.2 mA	2.4 mA	6.4 mA
<b>Application</b>	1.0 mA	0.7 mA	2.4 mA	3.2 mA



**Figure 4.22:** Battery duration of the target applications including the static power of the ADC with an increasing number of channels.

When comparing this work with the SoA solutions in Table 4.14, it is noteworthy that while some results and characteristics are similar, BioWolf excels in both intrusiveness and energy efficiency. For instance, the power consumption was improved by 8x while substantially reducing the output latency (<10ms vs. 300ms to [135]). Although the accuracy provided was similar, the HDC offered the chance of online and in-situ training (SVM requires offline analysis). This feature demonstrates that the system is ready for fast and intuitive usage, much required in everyday scenarios. Nevertheless, in this implementation, the advantages of Mr. Wolf SoC have not been fully exploited. Mr. Wolf has still more than 90% of CPU resources available for more advanced processing. We pushed these limits by increasing the number of EMG channels, targeting the new trend in highly dense EMG processing.

Fig. 4.22 shows the autonomy of the target architectures including the static power of the ADC and considering a battery of 100mAh. With a classification latency of 50ms, the STM32 and Mr. Wolf (1-core without built-ins) can only be scaled up to 64 channels because they are not able to meet this latency constraint.

**Table 4.14:** Comparison with SoA systems for EMG classification.

	Sensor	Processing	Intrusive	Accuracy	Harvester	Bat. Life
McIntosh [134] :	EMG Pressure	PC Offline	Yes	82	No	Nd
Huang[60] :	EMG IMU	PC Online	No	92	No	Nd
Liu [60] :	EMG	PC Online	No	80-90	No	24hs
Kartsch [60] :	EMG	Onboard Online	No	94	Yes	512hs/perp.
<b>This work :</b>	<b>EMG</b>	<b>Onboard Online</b>	<b>No</b>	<b>98-84</b>	<b>Yes</b>	<b>18/35hs<sup>1</sup></b>

<sup>1</sup> When scaling the system to achieve the same output rate as [135], and using a 60mAh battery, BioWolf can deliver up to 675hs of operation..

Nevertheless, Mr. Wolf denotes a more efficient operation by providing up to 3.9x more autonomy than the STM32. If we include the built-ins, the same core can now be scaled up 128 channels, providing more than 48h of continuous operation and demonstrating the benefits of the built-in instructions. Nonetheless, for the current application, Mr. Wolf (8-cores with built-ins) exceeds all other architectures (for every number of channels) mainly thanks to the distribution of the task into the different cores.

Another interesting fact comes from the line shapes, where all Mr. Wolf architectures will show a pseudo-constant average autonomy when increasing the number of channel up to 64 (with respect to four channels), with only 32% of autonomy degradation, while the commercial MCU will show peak degradation of 74%, i.e., a linear and steep decrease. This difference is due to the negligible power contribution of Mr. Wolf with respect to the ADC power. When scaling up the system with more channel, Mr. Wolf will still provide 19h to 12h of operation for 128 and 256 respectively, adequate for daily use. These results place Mr. Wolf as a unique enabler for the new generation of embedded systems.

## Chapter 5

# Conclusions

Human-Machine Interaction systems have recently reached a broader set of scenarios, including industrial, gaming, learning, and health tracking, thanks to advancements in Digital Signal Processing (DSP) and Machine Learning (ML) techniques. Porting applications to wearable and embedded systems, however, poses several challenges in terms of computational power, battery life, and wearability.

This dissertation addresses these challenges by providing complete solutions to reduce the intrusiveness of the system through sensor design and miniaturization of the hardware implementation while providing energy efficiency and real-time processing.

Significant effort has been granted to the develop of BioWolf, a multicore highly configurable platform that meets most of the modern HMI requirements for everyday scenarios. Intrusiveness has been significantly reduced through improvements in the sensor interface and PCB miniaturization, obtaining an area reduction of 4x to SoA systems. BioWolf can also deliver 4x more battery life than single-core commercial base platforms while only employing a fraction of the available processing power (<20%) and achieving SoA accuracy for both EEG and EMG applications (>90% for most cases). As BioWolf also provides energy harvesting techniques, some of the presented systems achieved self-sustainable operation or a significant boost in battery life (2-4x). This work highlighted the benefits of the multicore processing introduced by the PULP architecture in several biomedical study cases.

First, a sensor fusion approach for drowsiness detection was implemented. The system initially implemented a low-power ARM Cortex-M4 microcontroller, is capable of detecting five different levels of drowsiness with an average accuracy of 95.2% and a battery life of 6 hours, using a 200mAh battery. Energy optimizations achievable by accelerating the sensor fusion-based drowsiness detector on the PULP platform have been

also studied and motivated an actual implementation on BioWolf, resulting in a system offering an accuracy compliant with the SoA while providing significant advantages in terms of power consumption, allowing for 17x improvements in terms of power envelope.

In a second application a BCI featuring Canonical Correlation Analysis (CCA) of steady-state visual evoked potentials was implemented. The system achieved an average information transfer rate of 1.46 bits per second (bps) at a power budget of 6.31 mW, providing up to 38 hrs operation (65 mAh battery). To the best of the our knowledge, this design is the first to explore the significant energy boost of a parallel MCU to single-core MCUs for CCA-based BCI.

In a third application, BioWolf was employed for sensory evaluation through ERP measurements. The system provided online and real-time differentiation of two classes of images based only on EEG signals. By virtue of the energy efficiency of Mr. Wolf, the system exceeds 19h of battery life with a tiny 60 mAh LiPo battery. The preliminary results presented validate our embedded deployment and represent a step forward towards an unbiased automated food quality grading.

In the last application, an online approach for learning and classification of EMG-Based Gestures was presented. The system, based on the Hyperdimensional Computing and implemented in BioWolf can run a pattern recognition algorithm, recognizing up to 11 hand gestures, and ensure up to 18 h of continuous operation that can be further extended up to 35 h with outdoor illumination. This implementation once again validated the capabilities of BioWolf.

These achievements demonstrate that advancements in sensor interface and system miniaturization can be coupled with multicore architectures, such as PULP, to give growth to the next generation of unobtrusive and real-time embedded applications for biosignal processing. Still, the systems presented in this dissertation can benefit from further improvements regarding the accuracy, the number of features, sensor fusion, and advanced classification techniques, as currently, the resources available, especially in terms of computational power, are only partially employed.

# Bibliography

- [1] H. Ando, K. Takizawa, T. Yoshida, K. Matsushita, M. Hirata, and T. Suzuki. Wireless multichannel neural recording with a 128-mbps uwb transmitter for an implantable brain-machine interfaces. *IEEE Transactions on Biomedical Circuits and Systems*, 10(6):1068–1078, Dec 2016. ISSN 1932-4545. doi: 10.1109/TBCAS.2016.2514522.
- [2] X. Liu, M. Zhang, B. Subei, A. G. Richardson, T. H. Lucas, and J. Van der Spiegel. The pennbmbi: Design of a general purpose wireless brain-machine-brain interface system. *IEEE Transactions on Biomedical Circuits and Systems*, 9(2): 248–258, April 2015. ISSN 1932-4545. doi: 10.1109/TBCAS.2015.2392555.
- [3] Emotiv, 2018. URL <https://www.emotiv.com/>.
- [4] Neurosky. Neurosky MindWave, 2018. URL <http://neurosky.com/biosensors/eeg-sensor/biosensors/>.
- [5] OpenBCI, 2018. URL <http://openbci.com/>.
- [6] Gernot R Muller-Putz and Gert Pfurtscheller. Control of an electrical prosthesis with an SSVEP-based BCI. *IEEE Transactions on Biomedical Engineering*, 55(1): 361–364, 2008.
- [7] T Scott Saponas, Desney S Tan, Dan Morris, Ravin Balakrishnan, Jim Turner, and James A Landay. Enabling always-available input with muscle-computer interfaces. In *Proceedings of the 22nd annual ACM symposium on User interface software and technology*, pages 167–176. ACM, 2009.
- [8] touch bionics. <http://www.touchbionics.com/products>, 2018.
- [9] Ottobock. <https://www.ottobockus.com/prosthetics/upper-limb-prosthetics/solution-overview/myoelectric-prosthetics/>, 2018.
- [10] Marco Guermandi, Simone Benatti, Victor Javier Kartsch Morinigo, and Luca Bertini. A wearable device for minimally-invasive behind-the-ear eeg and evoked

- potentials. In *2018 IEEE Biomedical Circuits and Systems Conference (BioCAS)*, pages 1–4. IEEE, 2018.
- [11] Matthieu Duvinage, Thierry Castermans, Thierry Dutoit, M Petieau, T Hoellinger, C De Saedeleer, K Seetharaman, and G Cheron. A p300-based quantitative comparison between the emotiv epoc headset and a medical eeg device. *Biomedical Engineering*, 765(1):2012–2764, 2012.
- [12] Daniela De Venuto, Valerio F Annese, and Giovanni Mezzina. An embedded system remotely driving mechanical devices by p300 brain activity. In *Proceedings of the Conference on Design, Automation & Test in Europe*, pages 1014–1019. European Design and Automation Association, 2017.
- [13] DC Preston and BE Shapiro. Basic electromyography: analysis of motor unit action potentials. *Electromyography and neuromuscular disorders: clinical-electrophysiologic correlations*. Philadelphia: Elsevier, pages 215–29, 2005.
- [14] Lauren H Smith and Levi J Hargrove. Comparison of surface and intramuscular emg pattern recognition for simultaneous wrist/hand motion classification. In *2013 35th annual international conference of the IEEE engineering in medicine and biology society (EMBC)*, pages 4223–4226. IEEE, 2013.
- [15] Jeffrey W Britton, Lauren C Frey, JL Hopp, P Korb, MZ Koubeissi, WE Lievens, EM Pestana-Knight, and EK Louis St. *Electroencephalography (EEG): An introductory text and atlas of normal and abnormal findings in adults, children, and infants*. American Epilepsy Society, Chicago, 2016.
- [16] P Flandrin. Some aspects of non-stationary signal processing with emphasis on time-frequency and time-scale methods. In *Wavelets*, pages 68–98. Springer, 1989.
- [17] Saeid Sanei and Jonathon A Chambers. *Eeg signal processing*. 2007.
- [18] Mircea Steriade. 1. abstract 2. neuronal substrates of brain disconnection during nrem sleep 3. three types of brain rhythms during nrem sleep: the unified corticothalamic system 3.1. spindles, a thalamic rhythm under neocortical influence 3.2. delta waves: two different (thalamic and cortical) components 3.3. the neocortical slow oscillation groups thalamically generated nrem sleep rhythms. *Front Biosci*, 8:d878–99, 2003.
- [19] György Buzsáki. Theta rhythm of navigation: link between path integration and landmark navigation, episodic and semantic memory. *Hippocampus*, 15(7):827–840, 2005.



- [20] Jose L Cantero, Mercedes Atienza, Robert Stickgold, Michael J Kahana, Joseph R Madsen, and Bernat Kocsis. Sleep-dependent  $\theta$  oscillations in the human hippocampus and neocortex. *Journal of Neuroscience*, 23(34):10897–10903, 2003.
- [21] Jim Lagopoulos, Jian Xu, Inge Rasmussen, Alexandra Vik, Gin S Malhi, Carl F Eliassen, Ingrid E Arntsen, Jardar G Sæther, Stig Hollup, Are Holen, et al. Increased theta and alpha eeg activity during nondirective meditation. *The Journal of Alternative and Complementary Medicine*, 15(11):1187–1192, 2009.
- [22] Anat Perry, Shlomo Bentin, Idan Shalev, Salomon Israel, Florina Uzefovsky, Dori Bar-On, and Richard P Ebstein. Intranasal oxytocin modulates eeg mu/alpha and beta rhythms during perception of biological motion. *Psychoneuroendocrinology*, 35(10):1446–1453, 2010.
- [23] Dennis J McFarland, Laurie A Miner, Theresa M Vaughan, and Jonathan R Wolpaw. Mu and beta rhythm topographies during motor imagery and actual movements. *Brain topography*, 12(3):177–186, 2000.
- [24] Hanbing Lu, Yantao Zuo, Hong Gu, James A Waltz, Wang Zhan, Clara A Scholl, William Rea, Yihong Yang, and Elliot A Stein. Synchronized delta oscillations correlate with the resting-state functional mri signal. *Proceedings of the National Academy of Sciences*, 104(46):18265–18269, 2007.
- [25] Mathias Benedek, Sabine Bergner, Tanja Könen, Andreas Fink, and Aljoscha C Neubauer. Eeg alpha synchronization is related to top-down processing in convergent and divergent thinking. *Neuropsychologia*, 49(12):3505–3511, 2011.
- [26] John Polich and Jody Corey-Bloom. Alzheimer’s disease and p300: review and evaluation of task and modality. *Current Alzheimer Research*, 2(5):515–525, 2005.
- [27] Masaki Nakanishi et al. Enhancing Detection of SSVEPs for a high-speed brain speller using task-related component analysis. *IEEE Transactions on Biomedical Engineering*, 65(1), 2018.
- [28] Xiaogang Chen et al. High-speed spelling with a noninvasive brain–computer interface. *Proceedings of the national academy of sciences*, 112(44):E6058–E6067, 2015.
- [29] Xiaogang Chen, Zhikai Chen, Shangkai Gao, and Xiaorong Gao. A high-ITR SSVEP-based BCI speller. *Brain-Computer Interfaces*, 1(3-4):181–191, 2014.
- [30] Hsiang-Chih Chang et al. Real-time control of an SSVEP-actuated remote-controlled car. In *SICE Annual Conference 2010, Proceedings of*, pages 1884–1887. IEEE, 2010.

- [31] Zhonglin Lin et al. Frequency recognition based on canonical correlation analysis for SSVEP-based BCIs. *IEEE transactions on biomedical engineering*, 53(12):2610–2614, 2006.
- [32] Marc R Nuwer, Giancarlo Comi, Ronald Emerson, Anders Fuglsang-Frederiksen, Jean-Michel Guérit, Hermann Hinrichs, Akio Ikeda, Fransisco Jose C Luccas, and Peter Rappelsburger. Ifcn standards for digital recording of clinical eeg. *Electroencephalography and clinical Neurophysiology*, 106(3):259–261, 1998.
- [33] Marco Guermandi, Eleonora Franchi Scarselli, and Roberto Guerrieri. A driving right leg circuit (dgrl) for improved common mode rejection in bio-potential acquisition systems. *IEEE transactions on biomedical circuits and systems*, 10(2):507–517, 2016.
- [34] Marco Guermandi, Roberto Cardu, Eleonora Franchi Scarselli, and Roberto Guerrieri. Active electrode ic for eeg and electrical impedance tomography with continuous monitoring of contact impedance. *IEEE transactions on biomedical circuits and systems*, 9(1):21–33, 2015.
- [35] Jiawei Xu, Srinjoy Mitra, Akinori Matsumoto, Shrishail Patki, Chris Van Hoof, Kofi AA Makinwa, and Refet Firat Yazicioglu. A wearable 8-channel active-electrode eeg/eti acquisition system for body area networks. *IEEE Journal of Solid-State Circuits*, 49(9):2005–2016, 2014.
- [36] Alexander von Lühmann, Heidrun Wabnitz, Tilmann Sander, and Klaus-Robert Müller. M3ba: a mobile, modular, multimodal biosignal acquisition architecture for miniaturized eeg-nirs-based hybrid bci and monitoring. *IEEE Transactions on Biomedical Engineering*, 64(6):1199–1210, 2016.
- [37] Fan Zhang, Jeremy Holleman, and Brian P Otis. Design of ultra-low power biopotential amplifiers for biosignal acquisition applications. *IEEE transactions on biomedical circuits and systems*, 6(4):344–355, 2012.
- [38] Ultra-Low Power, Single-Channel Integrated Biopotential AFE. <https://datasheets.maximintegrated.com/en/ds/MAX30003.pdf>, 2019.
- [39] Electroencephalogram (EEG) Measurement - Analog Devices. <https://www.analog.com/en/applications/markets/healthcare-pavilion-home/vital-signs-measurement/eeg-measurement.html>, 2019.
- [40] Simone Benatti et al. Multiple biopotentials acquisition system for wearable applications. In *BIODEVICES*, pages 260–268, 2015.

- [41] Mattia Salvaro, Victor Kartsch, Simone Benatti, Michela Milano, and Luca Benini. Towards a novel hmi paradigm based on mixed eeg and indoor localization platforms. In *CAS (NGCAS), 2017 New Generation of*, pages 217–220. IEEE, 2017.
- [42] Matteo Rossi, Simone Benatti, Elisabetta Farella, and Luca Benini. Hybrid emg classifier based on hmm and svm for hand gesture recognition in prosthetics. In *Industrial Technology (ICIT), 2015 IEEE International Conference on*, pages 1700–1705. IEEE, 2015.
- [43] Xilin Liu, Jacob Sacks, Milin Zhang, Andrew G Richardson, Timothy H Lucas, and Jan Van der Spiegel. The virtual trackpad: An electromyography-based, wireless, real-time, low-power, embedded hand-gesture-recognition system using an event-driven artificial neural network. *IEEE Trans. Circuits Syst. II Express Briefs*, 64: 1257–1261, 2017.
- [44] D. Rossi. Sub-pj per operation scalable computing: The pulp experience. In *2016 IEEE SOI-3D-Subthreshold Microelectronics Technology Unified Conference (S3S)*, pages 1–3, Oct 2016. doi: 10.1109/S3S.2016.7804389.
- [45] Michael Gautschi, Davide Rossi, and Luca Benini. Customizing an open source processor to fit in an ultra-low power cluster with a shared l1 memory. In *Proceedings of the 24th Edition of the Great Lakes Symposium on VLSI, GLSVLSI '14*, pages 87–88, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2816-6. doi: 10.1145/2591513.2591569. URL <http://doi.acm.org/10.1145/2591513.2591569>.
- [46] Igor Loi, Davide Rossi, Germain Haugou, Michael Gautschi, and Luca Benini. Exploring multi-banked shared-l1 program cache on ultra-low power, tightly coupled processor clusters. In *Proceedings of the 12th ACM International Conference on Computing Frontiers, CF '15*, pages 64:1–64:8, New York, NY, USA, 2015. ACM. ISBN 978-1-4503-3358-0. doi: 10.1145/2742854.2747288. URL <http://doi.acm.org/10.1145/2742854.2747288>.
- [47] P. Davide Schiavone et al. Slow and steady wins the race? a comparison of ultra-low-power risc-v cores for internet-of-things applications. In *PATMOS*, pages 1–8, Sept 2017. doi: 10.1109/PATMOS.2017.8106976.
- [48] Simone Benatti, Fabio Montagna, Davide Rossi, and Luca Benini. Scalable eeg seizure detection on an ultra low power multi-core architecture. In *2016 IEEE Biomedical Circuits and Systems Conference (BioCAS)*, pages 86–89. IEEE, 2016.
- [49] Jose Antonio Urigüen and Begoña Garcia-Zapirain. Eeg artifact removal—state-of-the-art and guidelines. *Journal of neural engineering*, 12(3):031001, 2015.

- [50] Angkoon Phinyomark, Chusak Limsakul, and Pornchai Phukpattaranont. A novel feature extraction for robust emg pattern recognition. *arXiv preprint arXiv:0912.3973*, 2009.
- [51] Simone Benatti, Filippo Casamassima, Bojan Milosevic, Elisabetta Farella, Philipp Schönle, Schekeb Fateh, Thomas Burger, Qiuting Huang, and Luca Benini. A versatile embedded platform for emg acquisition and gesture recognition. *IEEE transactions on biomedical circuits and systems*, 9(5):620–630, 2015.
- [52] Zirui Lan, Olga Sourina, Lipo Wang, and Yisi Liu. Real-time eeg-based emotion monitoring using stable features. *The Visual Computer*, 32(3):347–358, 2016.
- [53] Amjed S Al-Fahoum and Ausilah A Al-Fraihat. Methods of eeg signal features extraction using linear analysis in frequency and time-frequency domains. *ISRN neuroscience*, 2014, 2014.
- [54] Simone Benatti, Fabio Montagna, Victor Kartsch, Abbas Rahimi, Davide Rossi, and Luca Benini. Online learning and classification of emg-based gestures on a parallel ultra-low power platform using hyperdimensional computing. *IEEE transactions on biomedical circuits and systems*, 13(3):516–528, 2019.
- [55] Texas Instruments, 2015. URL <http://www.ti.com/lit/ds/symlink/ads1298.pdf>.
- [56] A. Pullini, D. Rossi, I. Loi, A. Di Mauro, and L. Benini. Mr. wolf: A 1 gflop/s energy-proportional parallel ultra low power soc for iot edge processing. In *ESSCIRC 2018 - IEEE 44th European Solid State Circuits Conference (ESSCIRC)*, pages 274–277, Sept 2018. doi: 10.1109/ESSCIRC.2018.8494247.
- [57] Rifai Chai, Ganesh R Naik, Sai Ho Ling, and Hung T Nguyen. Hybrid brain–computer interface for biomedical cyber-physical system application using wireless embedded eeg systems. *Biomedical engineering online*, 16(1):5, 2017.
- [58] Victor Javier Kartsch et al. A sensor fusion approach for drowsiness detection in wearable ultra-low-power systems. *Information Fusion*, 43:66–76, 2018.
- [59] K. A. Condori, E. C. Urquizo, and D. A. Diaz. Embedded brain machine interface based on motor imagery paradigm to control prosthetic hand. In *2016 IEEE ANDESCON*, pages 1–4, Oct 2016. doi: 10.1109/ANDESCON.2016.7836266.
- [60] Jzau-Sheng Lin and Cheng-Hung Hsieh. A wireless BCI-controlled integration system in smart living space for patients. *Wireless Personal Communications*, 88(2):395–412, 2016.

- [61] Gang Li and Wan-Young Chung. Combined eeg-gyroscope-tdcs brain machine interface system for early management of driver drowsiness. *IEEE Transactions on Human-Machine Systems*, 48(1):50–62, 2018.
- [62] Yu Mike Chi et al. Dry and noncontact EEG sensors for mobile brain–computer interfaces. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 20(2):228–235, 2012.
- [63] C. M. McCrimmon, J. L. Fu, M. Wang, L. S. Lopes, P. T. Wang, A. Karimi-Bidhendi, C. Y. Liu, P. Heydari, Z. Nenadic, and A. H. Do. Performance assessment of a custom, portable, and low-cost brain–computer interface platform. *IEEE Transactions on Biomedical Engineering*, 64(10):2313–2320, Oct 2017. ISSN 0018-9294. doi: 10.1109/TBME.2017.2667579.
- [64] Tapas Kumar Aich. Absent posterior alpha rhythm: An indirect indicator of seizure disorder? *Indian journal of psychiatry*, 56(1):61, 2014.
- [65] D.F. Dinges, M. M. Mallis, G. Maislin, and J. W. Powell IV. Evaluation of techniques for ocularmeasurement as an index of fatigue andthe basis for alertness management. In *National Highway Traffic Safety Administration*, page 113, 1998.
- [66] Hua Cai and Yingzi Lin. A roadside its data bus prototype for intelligent highways. *IEEE Transactions on intelligent transportation systems*, 9(2):344–348, 2008.
- [67] Chuan Xu, Xuesong Wang, and Xiaohong Chen. *Modeling Fatigue Level by Driver’s Lane-Keeping Indicators*, pages 2282–2288. doi: 10.1061/9780784413159.331. URL <http://ascelibrary.org/doi/abs/10.1061/9780784413159.331>.
- [68] MarcoJavier Flores, JoséMaría Armingol, and Arturo de la Escalera. Driver drowsiness warning system using visual information for both diurnal and nocturnal illumination conditions. *EURASIP Journal on Advances in Signal Processing*, 2010(1):438205, 2010. doi: 10.1155/2010/438205. URL <http://dx.doi.org/10.1155/2010/438205>.
- [69] Qiang Ji, Zhiwei Zhu, and P. Lan. Real-time nonintrusive monitoring and prediction of driver fatigue. *IEEE Transactions on Vehicular Technology*, 53(4):1052–1068, July 2004. ISSN 0018-9545. doi: 10.1109/TVT.2004.830974.
- [70] A. Kolli, A. Fasih, F. A. Machot, and K. Kyamakya. Non-intrusive car driver’s emotion recognition using thermal camera. In *Proceedings of the Joint INDS’11 ISTET’11*, pages 1–5, July 2011. doi: 10.1109/INDS.2011.6024802.
- [71] G. Fortino, R. Giannantonio, R. Gravina, P. Kuryloski, and R. Jafari. Enabling effective programming and flexible management of efficient body sensor network applications. *IEEE Transactions on Human-Machine Systems*, 43(1):115–133, 2013.

- [72] Raffaele Gravina, Parastoo Alinia, Hassan Ghasemzadeh, and Giancarlo Fortino. Multi-sensor fusion in body sensor networks: State-of-the-art and research challenges. *Information Fusion*, 35:68 – 80, 2017. ISSN 1566-2535. doi: <http://dx.doi.org/10.1016/j.inffus.2016.09.005>. URL <http://www.sciencedirect.com/science/article/pii/S156625351630077X>.
- [73] Bahador Khaleghi, Alaa Khamis, Fakhreddine O. Karray, and Saiedeh N. Razavi. Multisensor data fusion: A review of the state-of-the-art. *Information Fusion*, 14(1):28 – 44, 2013. ISSN 1566-2535. doi: <http://dx.doi.org/10.1016/j.inffus.2011.08.001>. URL <http://www.sciencedirect.com/science/article/pii/S1566253511000558>.
- [74] Giancarlo Fortino, Stefano Galzarano, Raffaele Gravina, and Wenfeng Li. A framework for collaborative computing and multi-sensor data fusion in body sensor networks. *Information Fusion*, 22:50 – 70, 2015. ISSN 1566-2535. doi: <http://dx.doi.org/10.1016/j.inffus.2014.03.005>. URL <http://www.sciencedirect.com/science/article/pii/S156625351400044X>.
- [75] Amardeep Sathyanarayana, Pinar Boyraz, and John H.L. Hansen. Information fusion for robust  $\hat{\sim}$  context and driver aware<sup>TM</sup> active vehicle safety systems. *Information Fusion*, 12(4):293 – 303, 2011. ISSN 1566-2535. doi: <http://dx.doi.org/10.1016/j.inffus.2010.06.004>. URL <http://www.sciencedirect.com/science/article/pii/S1566253510000679>. Special Issue on Information Fusion for Cognitive Automobiles.
- [76] Boon-Giin Lee and Wan-Young Chung. A smartphone-based driver safety monitoring system using data fusion. *Sensors*, 12(12):17536, 2012. ISSN 1424-8220. doi: 10.3390/s121217536. URL <http://www.mdpi.com/1424-8220/12/12/17536>.
- [77] Angelica Reyes-Muñoz, Mari Carmen Domingo, Marco Antonio López-Trinidad, and José Luis Delgado. Integration of body sensor networks and vehicular ad-hoc networks for traffic safety. *Sensors*, 16(1):107, 2016. ISSN 1424-8220. doi: 10.3390/s16010107. URL <http://www.mdpi.com/1424-8220/16/1/107>.
- [78] Chin-Teng Lin, Ruei-Cheng Wu, Sheng-Fu Liang, Wen-Hung Chao, Yu-Jie Chen, and Tzyy-Ping Jung. Eeg-based drowsiness estimation for safety driving using independent component analysis. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 52(12):2726–2738, Dec 2005. ISSN 1549-8328.
- [79] C. T. Lin, C. J. Chang, B. S. Lin, S. H. Hung, C. F. Chao, and I. J. Wang. A real-time wireless brain-computer interface system for drowsiness detection. *IEEE Transactions on Biomedical Circuits and Systems*, 4(4):214–222, Aug 2010. ISSN 1932-4545.

- [80] Unsoo Ha and Hoi-Jun Yoo. A multimodal drowsiness monitoring ear-module system with closed-loop real-time alarm. In *IEEE Biomedical Circuits and Systems Conference (BioCAS)*, Oct 2016.
- [81] P. Li, R. Meziane, M. J. D. Otis, H. Ezzaidi, and P. Cardou. A smart safety helmet using imu and eeg sensors for worker fatigue detection. In *2014 IEEE International Symposium on Robotic and Sensors Environments (ROSE) Proceedings*, pages 55–60, Oct 2014.
- [82] Gianluca Borghini, Laura Astolfi, Giovanni Vecchiato, Donatella Mattia, and Fabio Babiloni. Measuring neurophysiological signals in aircraft pilots and car drivers for the assessment of mental workload, fatigue and drowsiness. *Neuroscience Biobehavioral Reviews*, 44:58 – 75, 2014. ISSN 0149-7634. doi: <http://dx.doi.org/10.1016/j.neubiorev.2012.10.003>. URL <http://www.sciencedirect.com/science/article/pii/S0149763412001704>. Applied Neuroscience: Models, methods, theories, reviews. A Society of Applied Neuroscience (SAN) special issue.
- [83] Robert Schleicher, Niels Galley, Susanne Briest, and Lars Galley. Blinks and saccades as indicators of fatigue in sleepiness warnings: looking tired? *Ergonomics*, 51(7):982–1010, 2008.
- [84] MW Johns et al. The amplitude-velocity ratio of blinks: a new method for monitoring drowsiness. *Sleep*, 26:A51, 2003.
- [85] A. Picot et al. On-line automatic detection of driver drowsiness using a single electroencephalographic channel. In *EMBC*, Aug 2008.
- [86] J. Park et al. Wireless dry eeg for drowsiness detection. In *EMBC*, pages 3298–3301, Aug 2011.
- [87] Ward Vanlaar, Herb Simpson, Dan Mayhew, and Robyn Robertson. Fatigued and drowsy driving: A survey of attitudes, opinions and behaviors. *Journal of Safety Research*, 39(3):303 – 309, 2008. ISSN 0022-4375. doi: <http://dx.doi.org/10.1016/j.jsr.2007.12.007>. URL <http://www.sciencedirect.com/science/article/pii/S0022437508000492>.
- [88] Anneke Heitmann, Rainer Guttkuhn, Acacia Aguirre, Udo Trutschel, and Martin Moore-Ede. Technologies for the monitoring and prevention of driver fatigue. In *Proceedings of the first international driving symposium on human factors in driver assessment, training and vehicle design*, volume 86, 2001.

- [89] Mohamad hoseyn Sigari, Muhammad reza Pourshahabi, Mohsen Soryani, and Mahmood Fathy. A review on driver face monitoring systems for fatigue and distraction detection.
- [90] A Muzet, T Pébayle, J Langrognet, and S Otmani. Awake pilot study no. 2: Testing steering grip sensor measures. *CEPA, Gatineau, QC, Canada, Tech. Rep. IST-2000-28062*, 2003.
- [91] A Anund et al. Pilot 15—report: Vti. *Swedish National Road and Transport Research Institute*, 2004.
- [92] I.V.D.O. William O. Tatum. *Handbook of EEG Interpretation, Second Edition*. Springer Publishing Company, 2014. ISBN 9781617051807. URL <https://books.google.it/books?id=PZnSCgAAQBAJ>.
- [93] ARM. <https://www.arm.com/products/processors/cortex-m/cortex-microcontroller-software-interface-standard.php>, 2016.
- [94] Mark E. Howard, Melinda L. Jackson, David Berlowitz, Fergal O’Donoghue, Philip Swann, Justine Westlake, Vanessa Wilkinson, and Rob J. Pierce. Specific sleepiness symptoms are indicators of performance impairment during sleep deprivation. *Accident Analysis Prevention*, 62:1 – 8, 2014. ISSN 0001-4575. doi: <http://dx.doi.org/10.1016/j.aap.2013.09.003>. URL <http://www.sciencedirect.com/science/article/pii/S0001457513003564>.
- [95] Arun Sahayadhas et al. Detecting driver drowsiness based on sensors: a review. *Sensors*, 12(12):16937–16953, 2012.
- [96] Qiang Ji and Xiaojie Yang. Real-time eye, gaze, and face pose tracking for monitoring driver vigilance. *Real-Time Imaging*, 8(5):357 – 377, 2002. ISSN 1077-2014. doi: <http://dx.doi.org/10.1006/rtim.2002.0279>. URL <http://www.sciencedirect.com/science/article/pii/S1077201402902792>.
- [97] Amna Rahman et al. Real time drowsiness detection using eye blink monitoring. In *NSEC*, pages 1–7. IEEE, 2015.
- [98] Valer Jurcak, Daisuke Tsuzuki, and Ipeita Dan. 10/20, 10/10, and 10/5 systems revisited: their validity as relative head-surface-based positioning systems. *Neuroimage*, 34(4):1600–1611, 2007.
- [99] Fabio Montagna et al. Flexible, scalable and energy efficient bio-signals processing on the pulp platform: A case study on seizure detection. *Journal of Low Power Electronics and Appl.*, 2017.



- [100] Hong J. Eoh, Min K. Chung, and Seong-Han Kim. Electroencephalographic study of drowsiness in simulated driving with sleep deprivation. *International Journal of Industrial Ergonomics*, 2005. ISSN 0169-8141. doi: <https://doi.org/10.1016/j.ergon.2004.09.006>.
- [101] Victor Kartsch et al. A wearable eeg-based drowsiness detection system with blink duration and alpha waves analysis. In *NER*, 2017.
- [102] Sugato Basu, Arindam Banerjee, and Raymond Mooney. Semi-supervised clustering by seeding. In *In Proceedings of 19th International Conference on Machine Learning (ICML-2002*. Citeseer, 2002.
- [103] Victor Kartsch, Giuseppe Tagliavini, Marco Guermandi, Simone Benatti, Davide Rossi, and Luca Benini. Biowolf: A sub-10 mw 8-channel advanced brain computer interface platform with a 9-core processor and ble connectivity. *IEEE transactions on biomedical circuits and systems*, 2019.
- [104] Lawrence Ashley Farwell and Emanuel Donchin. Talking off the top of your head: toward a mental prosthesis utilizing event-related brain potentials. *Electroencephalography and clinical Neurophysiology*, 70(6):510–523, 1988.
- [105] Emanuel Donchin et al. The mental prosthesis: assessing the speed of a P300-based brain-computer interface. *IEEE transactions on rehabilitation engineering*, 8(2):174–179, 2000.
- [106] Hilit Serby et al. An improved P300-based brain-computer interface. *IEEE Transactions on neural systems and rehabilitation engineering*, 13(1):89–98, 2005.
- [107] William M Perlstein et al. Steady-state visual evoked potentials reveal frontally-mediated working memory activity in humans. *Neuroscience letters*, 342(3):191–195, 2003.
- [108] Ming Cheng et al. Design and implementation of a brain-computer interface with high transfer rates. *IEEE transactions on biomedical engineering*, 49(10):1181–1186, 2002.
- [109] Niccolò Mora et al. Plug&play brain-computer interfaces for effective active and assisted living control. *Medical & biological engineering & computing*, 55(8):1339–1352, 2017.
- [110] J Vernon Odom et al. Visual evoked potentials standard (2004). *Documenta ophthalmologica*, 108(2):115–123, 2004.

- [111] Nikolay V Manyakov et al. Sampled sinusoidal stimulation profile and multichannel fuzzy logic classification for monitor-based phase-coded ssvep brain-computer interfacing. *Journal of neural engineering*, 10(3):036011, 2013.
- [112] Zahra M. Aghajan, Peter Schuette, Tony A. Fields, Michelle E. Tran, Sameed M. Siddiqui, Nicholas R. Hasulak, Thomas K. Tcheng, Dawn Eliashiv, Emily A. Mankin, John Stern, Itzhak Fried, and Nanthia Suthana. Theta oscillations in the human medial temporal lobe during real-world ambulatory movement. *Current Biology*, 27(24):3743 – 3751.e3, 2017. ISSN 0960-9822. doi: <https://doi.org/10.1016/j.cub.2017.10.062>. URL <http://www.sciencedirect.com/science/article/pii/S0960982217313994>.
- [113] Gene H Golub. Matrix decompositions and statistical calculations. In *Statistical Computation*, pages 365–397. Elsevier, 1969.
- [114] M. Salvaro and S. Benatti and V. Kartsch and M. Guermandi and L. Benini. A Minimally Invasive Low-Power Platform for Real-Time Brain Computer Interaction based on Canonical Correlation Analysis. *IEEE Internet of Things Journal*, pages 1–1, 2018. ISSN 2327-4662. doi: 10.1109/JIOT.2018.2866341.
- [115] Pablo F Diez et al. Asynchronous BCI control using high-frequency SSVEP. *Journal of neuroengineering and rehabilitation*, 8(1):39, 2011.
- [116] Martin Spüler. A high-speed brain-computer interface (BCI) using dry EEG electrodes. *PloS one*, 12(2), 2017.
- [117] Hubert Cecotti. A self-paced and calibration-less SSVEP-based brain-computer interface speller. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 18(2):127–133, 2010.
- [118] Herbert Stone and Joel L. Sidel. Sensory evaluation practices (second edition). Food Science and Technology, pages 1 – 17. Academic Press, London, second edition edition, 1993. ISBN 978-0-12-672482-0. doi: <https://doi.org/10.1016/B978-0-12-672482-0.50008-3>.
- [119] International Organization for Standardization ISO. Sensory analysis methodology. general guidance for establishing a sensory profile. ISO 13299:2016.
- [120] Harald Schupp et al. Affective picture processing: the late positive potential is modulated by motivational relevance. *Psychophysiology*, 2000.
- [121] Stephen Brown, Henk van Steenbergen, Guido Band, Mischa de Rover, and Sander Nieuwenhuis. Functional significance of the emotion-related late positive potential. *Frontiers in Human Neuroscience*, 6:33, 2012. ISSN 1662-5161. doi:

- 10.3389/fnhum.2012.00033. URL <https://www.frontiersin.org/article/10.3389/fnhum.2012.00033>.
- [122] Thad Starner, Joshua Weaver, and Alex Pentland. Real-time american sign language recognition using desk and wearable computer based video. *IEEE Transactions on pattern analysis and machine intelligence*, 20(12):1371–1375, 1998.
- [123] Mohammadreza Asghari Oskoei, Huosheng Hu, et al. Support vector machine-based classification scheme for myoelectric control applied to upper limb. *IEEE Trans. Biomed. Engineering*, 55(8):1956–1965, 2008.
- [124] Manfredo Atzori, Matteo Cognolato, and Henning Müller. Deep learning with convolutional neural networks applied to electromyography data: A resource for the classification of movements for prosthetic hands. *Frontiers in neurorobotics*, 10:9, 2016.
- [125] Bojan Milosevic, Simone Benatti, and Elisabetta Farella. Design challenges for wearable emg applications. In *2017 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pages 1432–1437. IEEE, 2017.
- [126] Thalmic Labs. Thalmic’s MYO Armband. URL <https://www.myo.com/>.
- [127] Jun Liu, Fan Zhang, and He Helen Huang. An open and configurable embedded system for emg pattern recognition implementation for artificial arms. In *Engineering in Medicine and Biology Society (EMBC), 2014 36th Annual International Conference of the IEEE*, pages 4095–4098. IEEE, 2014.
- [128] Xiaorong Zhang, He Huang, and Qing Yang. Real-time implementation of a self-recovery emg pattern recognition interface for artificial arms. In *Engineering in Medicine and Biology Society (EMBC), 2013 35th Annual International Conference of the IEEE*, pages 5926–5929. IEEE, 2013.
- [129] Paolo Gentile, Marco Pessione, Antonio Suppa, Alessandro Zampogna, and Fernanda Irrera. Embedded wearable integrating real-time processing of electromyography signals. In *Multidisciplinary Digital Publishing Institute Proceedings*, volume 1, page 600, 2017.
- [130] Pentti Kanerva. Hyperdimensional computing: An introduction to computing in distributed representation with high-dimensional random vectors. *Cognitive Computation*, 1(2):139–159, 2009. ISSN 1866-9956. doi: 10.1007/s12559-009-9009-8. URL <http://dx.doi.org/10.1007/s12559-009-9009-8>.
- [131] Abbas Rahimi, Simone Benatti, Pentti Kanerva, Luca Benini, and Jan M. Rabaey. Hyperdimensional biosignal processing: A case study for EMG-based hand gesture

- recognition. In *IEEE International Conference on Rebooting Computing*, October 2016.
- [132] Fabio Montagna, Abbas Rahimi, Simone Benatti, Davide Rossi, and Luca Benini. Pulp-hd: Accelerating brain-inspired high-dimensional computing on a parallel ultra-low power platform. In *Proceedings of the 55th Annual Design Automation Conference, DAC '18*, pages 111:1–111:6, New York, NY, USA, 2018. ACM. ISBN 978-1-4503-5700-5. doi: 10.1145/3195970.3196096. URL <http://doi.acm.org/10.1145/3195970.3196096>.
- [133] S. Benatti, G. Rovere, J. Bösser, F. Montagna, E. Farella, H. Glaser, P. Schönle, T. Burger, S. Fateh, Q. Huang, and L. Benini. A sub-10mw real-time implementation for emg hand gesture recognition based on a multi-core biomedical soc. In *2017 7th IEEE International Workshop on Advances in Sensors and Interfaces (IWASI)*, pages 139–144, June 2017. doi: 10.1109/IWASI.2017.7974234.
- [134] Jess McIntosh, Charlie McNeill, Mike Fraser, Frederic Kerber, Markus Löchtefeld, and Antonio Krüger. Empress: Practical hand gesture classification with wrist-mounted emg and pressure sensing. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 2332–2342. ACM, 2016.
- [135] Victor Kartsch, Simone Benatti, Maurizio Mancini, Michele Magno, and Luca Benini. Smart wearable wristband for emg based gesture recognition powered by solar energy harvester. In *2018 IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 1–5. IEEE, 2018.