

Alma Mater Studiorum Università di Bologna
Archivio istituzionale della ricerca

An Energy-Efficient IoT node for HMI applications based on an ultra-low power Multicore Processor

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Kartsch, V., Guermandi, M., Benatti, S., Montagna, F., Benini, L. (2019). An Energy-Efficient IoT node for HMI applications based on an ultra-low power Multicore Processor. Institute of Electrical and Electronics Engineers Inc. [10.1109/SAS.2019.8705984].

Availability:

This version is available at: <https://hdl.handle.net/11585/708143> since: 2019-12-12

Published:

DOI: <http://doi.org/10.1109/SAS.2019.8705984>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

This is the post peer-review accepted manuscript of:

V. Kartsch, M. Guermandi, S. Benatti, F. Montagna and L. Benini, "An Energy-Efficient IoT node for HMI applications based on an ultra-low power Multicore Processor", 2019 IEEE Sensors Applications Symposium (SAS), Sophia Antipolis, France, 2019, pp. 1-6. doi: 10.1109/SAS.2019.8705984

The published version is available online at: <https://doi.org/10.1109/SAS.2019.8705984>

© 2019 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works

An Energy-Efficient IoT node for HMI applications based on an ultra-low power Multicore Processor

Victor Kartsch^{†‡}, Marco Guermandi^{†‡}, Simone Benatti[†], Fabio Montagna[†], Luca Benini^{†‡},

[†]DEI, University of Bologna, Italy.

[†]Email: {victorjavier.kartsch, marco.guermandi, simone.benatti, fabio.montagna, luca.benini}@unibo.it

[‡]Integrated System Laboratory, ETHZ, Zurich, Switzerland. Email: lbenini@iis.ee.ethz.ch

Abstract—Developing 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 paper we present a smart sensor node for IoT and HMI based on a programmable Parallel Ultra-Low-Power (PULP) platform. We tested the system on a hand gesture recognition application, which is a preferred way of interaction in HMI design. A wearable armband with 8 EMG sensors is controlled by our IoT node, running a machine learning algorithm in real-time, recognizing up to 11 gestures with a power envelope of 11.84 mW. As a result, the proposed approach is capable 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.

Keywords—Embedded systems, ultra-low power, multi-core, PULP, EMG.

I. INTRODUCTION

The global Human-Machine Interface (HMI) market is expected to generate revenues of more than 8 billion USD over the next 5 years. This trend is driven by the increasing adoption of devices for industrial automation [1], wearable health tracking [2] and, more in general, the growing plethora of IoT ecosystems.

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, including industrial control, healthcare, gaming and rehabilitation.

Decoding human intentions expressed by hand gestures is usually based on two main approaches: (i) visual recognition of hand gestures using computer vision techniques [3]; (ii) recognition based on the analysis of the electrical activity of the muscles involved in the gestures [4]. The former solution relies on the image processing of gesture captured by video cameras. Based on machine learning algorithms, it can recognize a large number of gestures [3], but it requires an external infrastructure such as fixed cameras, mounting attachments, and power supply, and it is very sensitive to environmental factors, such as variations of the intensity of lighting or line of sight interruption.

The alternative approach is based on decoding ElectroMyography (EMG) signal by leveraging techniques ranging from direct control [5] to pattern recognition [6], to deep learning

[7] and synergies [8], with the objective of mapping muscle contractions onto the corresponding hand gesture.

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 [9]. Some attempts have been made at a commercial level, such as the MYO [10], 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 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.

In an effort to move towards fully portable solutions, an approach which is gaining traction in the last year 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 acquisition and processing of EMG signals, which are 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 [11] or [12], 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 [13] and [14] 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 pose two significant challenges for the digital processing part: (i) the power envelope of the digital platforms must be minimized to allow high-quality signal acquisition via an Analog-Front-End (AFE) and (ii) approaches based on data streaming, which offloads the signal processing on external platforms, do not scale well because of limited bandwidth and high energy-per-transmitted bit of wireless interfaces, even though energy-efficient protocols are used (e.g. Bluetooth Low Energy). In this work, we introduce BioWolf, an integrated platform for computationally-intensive medical IoT applications, which addresses all these challenges as it

provides an ULP compute platform that can process biosignals in parallel and locally with a power budget lower than that of the AFE. Our platform is based on Mr. Wolf [15], a programmable Parallel Ultra-Low-Power processor that combines high versatility and compute efficiency higher than single-core architectures such as those available in standard MCUs with wireless connectivity. Hence, local end-to-end processing (i.e., with on-board classification) has also a lower power budget than streaming and remote recognition (in addition to lower latency and more robustness wrt wireless connectivity issues), employing 2.4x and 7x less power than the AFE and direct data streaming, respectively.

The PULP processor is coupled with a commercial Bluetooth Low Energy (BLE) SoC (Nordic nRF52832), which enables communications and auxiliary support for the system. The board also integrates an 8-channel Analog Front End (AFE) for the analog-to-digital conversion of the input signals. The system also includes an Energy Harvesting (EH) subsystem that provides extended battery life and automated battery recharging. 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. To validate the system, we integrated it in an elastic armband, to enable a hand gesture recognition device, based on Hyperdimensional Computing [16], a novel pattern recognition framework. First, we validate the electrical characteristic of the signal acquisition, demonstrating the suitability of Biowolf for biosignal processing, then 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 is further extended when energy is generated 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.

II. MATERIAL AND METHODS

A. Embedded Architecture

BioWolf is a highly-configurable platform for acquisition and embedded processing of biopotentials featuring a Parallel Ultra-Low-Power (PULP) SoC MCU for signal processing, an ARM-based Nordic SoC MCU for Bluetooth Low Energy (BLE) communications and system management, an Analog Front End (AFE) for analog-to-digital conversion of biosignals and a nano-power buck-boost regulator for energy harvesting. A T.I. BQ27441 fuel gauge is also present allowing to regularly check for battery status on a I2C interface. Fig. 1 shows a block diagram of the complete system and Fig. 2 shows the final PCB implementation.

Mr. Wolf, the Nordic SoC and the AFE are connected via SPI bus. Three operating modes are available, as described below.




Fig. 1. BioWolf System Architecture.




Fig. 2. BioWolf Board. Top side allocates Mr. Wolf, the AFE and part of the power supply section. Bottom side is mostly dedicated to the nRF52832 SoC, fuel gauge, connectors and the analog power supply section.

- When data needs to be streamed out directly (eventually after some basic processing such as simple filtering), Mr. Wolf is put in sleep mode and the Nordic SoC acts as master on the SPI bus, reading data from the AFE.
- When more computationally intensive processing is required, Mr. Wolf guarantees the best power efficiency to the system and is therefore the one controlling the SPI bus as the master, reading data from the AFE, processing it and sending only the result of such processing to the Nordic SoC for BLE transmission.
- When the system is not required to acquire and/or process data, it can be put in a deep sleep mode to minimize power consumption. Wake up is obtained by putting the device in a NFC field, such as tapping on it with a NFC-enabled smart-phone or tablet.

Biosignals are acquired by a multichannel commercial AFE from TI (ADS1298). The AFE is the *de-facto* standard used in biopotential acquisition platforms and presents a very favorable trade-off between performance and power consumption, since its 3 V single supply does not require step-up DC/DC conversion of the battery voltage, without significantly affecting noise performance. The board supports simultaneous sampling of up to 8 differential channels at frequencies up to 32 kbps with a gain of the input programmable gain amplifier (PGA) from 1 to 12 and a maximum resolution of 24-bits. The system is compatible both with dry and wet electrodes.

Mr. Wolf is a multi-core programmable SoC implemented in CMOS 40nm technology that combines a tiny (12 K gates)

RISC-V processor (zero-risky) [17], namely the Fabric Controller (FC), with a cluster of eight RISC-V processors equipped with flexible and powerful DSP extensions available on the RI5CY processor [17]. The cluster is coupled with a single-cycle latency multi-banked L1 memory (64 kB) allowing fast data transfer among the cores, and with an 'off the cluster' 512 kB of memory (L2) with 15 cycles latency. A dedicated DMA controller allows reducing the latency and computational power associated with data transfer. It also features two floating-point units (FPU) that are shared among the cores. Mr. Wolf can achieve very fine-grained parallelism and high energy efficiency in parallel workloads through a dedicated hardware block (HW Sync) that provides fast event management, parallel thread dispatching and synchronization. The SoC contains a full set of peripherals, including a Quad SPI (QSPI), I2C and UART, with data transfers also managed by a multi-channel I/O DMA to reduce the load on the system. In run mode, the SoC is powered by an internal DC/DC converter that can be programmed to deliver from 0.8 V to 1.1 V. In sleep mode, a low-dropout (LDO) regulator powers a real-time clock (32 kHz crystal oscillator) that controls a programmed wake-up and, optionally, part of the L2 memory, allowing retention of application state for fast wake-up. In deep sleep mode, the power consumption of the MCU is about 108 μ W that can be further reduced to 72 μ W when no retention is required.

Data communication (and basic processing if needed) is performed by the nRF52832 SoC from Nordic. The MCU, based on an ARM Cortex-M4 (up to 64 MHz clock frequency) provides flexible Bluetooth 5 (BLE) communication at a low-power budget. This MCU also serves as a device manager of the board. It allows choosing the operation mode (sleep, raw data streaming, data acquisition and processing), including programming Mr. Wolf accordingly and setting power on/down of the analog section. It also detects battery status from the fuel gauge.

Power supply, battery management, and energy harvesting are managed by a Texas Instruments BQ25570. The IC implements a Maximum Power Point Tracking (MPPT) that adapts the input impedance of the solar cells maximizing the energy conversion in all the lighting conditions with up to 90% of efficiency. This energy is then used to recharge a small factor 65 mAh LiPo battery. The Energy Subsystem (EH) also provides a high efficient buck converter that delivers a stable voltage output of 1.8 V to supply the digital portions of the board. An additional output is available, connected to the battery voltage when its voltage level is higher than 3 V. This is used to power the analog portions of the board, in particular, the AFE which requires a minimum supply voltage of 2.7 V.

B. Hyperdimensional Computing

To demonstrate the performance of our system architecture, we propose as a case study the classification of hand gestures from EMG signal through HD Computing algorithm, a brain-inspired approach that computes with points in the HD space (hypervectors) as an alternative to numbers [16].




Fig. 3. Implementation on BioWolf of the HD computing algorithm.

To exploit all the capabilities of the hardware implementation, these hypervectors are considered as (pseudo)random dense binary vectors composed of an equal number of randomly placed 0s and 1s, which can be combined into new hypervectors through well-defined algebraic operations such as componentwise XOR (\oplus) as multiplication, the componentwise majority function ($[+]$) as addition, and one-bit circular rotation (ρ) as permutation. Features are extracted from the raw signals and mapped (i.e. encoded) into the HD space using Item Memory (IM) and Continuous Item Memory (CIM) [18] matrices. 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. Discretizing the features in K levels, we have K hypervectors ($V_1 \dots V_K$) where V_1 and V_K are related to the minimum and maximum input values and the intermediate levels are generated by a linear interpolation between these two orthogonal endpoints [18]. The HD computing provides two encoders, spatial and temporal. The first one captures the spatial information contained in the signal with a componentwise XOR between E and V resulting (at instant t):

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

Sometimes the spatial information is not enough, and the temporal information is required. This can be done by a temporal encoder that extracts such information 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:

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

where ρ^k stands for k times permutation. The HD computing is trained off-line, generating different n -grams for each gesture and adding them to create a *prototype* hypervector stored in the associative memory (AM). During inference, an unseen feature is encoded into an n -gram (*query*) hypervector, compared with all the prototype hypervectors in AM through the Hamming distance. Thus, the label associated with the minimum distance is assigned as the classification output. Fig. 3 summarizes the classification process introduced above.

C. Implementation and Optimization on BioWolf

Typically, binary hypervectors assume a very high dimension (i.e., 10k-D), and they can be manipulated using multiplication, addition, and permutation (MAP) operations after compacting them into 32-bit unsigned integer, leading to a conspicuous gain in performance and memory requirements. This representation requires bitwise operations (i.e. read/insert bits




Fig. 4. 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.

into a 32-bit word) and to count the number of 1s in a word (the well-known popcount). The RI5CY processor allows aggressive performance optimizations including bit manipulation instructions (builtins). This allows bitwise operations in 1 clock cycle [19], dramatically reducing the computational load on the MCU. An other optimization derives from the exploiting of the parallel programming models through an optimized version of Open Multi-Processing (OpenMP).

III. EXPERIMENTAL RESULTS

A. Electrical characterization

We characterized the system at 1000 samples-per-second (SPS) sampling frequency, that guarantees a bandwidth of 262 Hz, exceeding the needs of most target applications. Noise is measured by shorting the inputs of the electrodes and varies depending on the chosen PGA gain. We compare the performance with IFCN standards for clinical recording of EEG signals [20], which are generally considered as the most stringent for bio-potential acquisition. With PGA gain equal to 1, it is measured at $1.65 \mu V_{RMS}$ in the 0.5-100 Hz band, decreasing to $0.97 \mu V_{RMS}$ (gain = 2), $0.49 \mu V_{RMS}$ (gain = 4) and $0.41 \mu V_{RMS}$ (gain = 12) with PGA gain equal to 12. Common Mode Rejection Ratio for a 50 Hz, 2 Vpp signal ranges from a minimum of 115 dB ($G = 1$) to 122 dB (gain = 12). Channel isolation exceeds 100 dB. These values are in line with IFCN standards for clinical recording of EEG signals.

We also estimated the harvesting capabilities of the system by measuring the current applied by the EH subsystem to the battery in different illuminations. The installed solar panel has the same footprint of the board (2 x 4 cm) aiming to preserve wearability. Figure 4 summarizes the harvester performances denoting, at the magnified frame, that under indoor illumination (≈ 600 lux), generated current is quite low (around $80 \mu A$) but still enough to charge the system when in standby (around $80 \mu A$ current consumption, as shown in subsection III-C). This situation dramatically improves when moving into brighter environments, where the solar panel can deliver up to 2.5 mA.




Fig. 5. Average accuracy obtaining by HD computing, using the same data collected by 10 subjects, increasing the number of gestures (from 1 to 11).

B. HDC performance

To demonstrate the performance of the system in terms of classification accuracy, we involved in the experiment ten able-bodied subjects (aged 26-42) without a previous history of neurological or muscular disorders. All participants provided written consent to participate in the experiments.

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 paper. 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. Fig. 5 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 [14].

Table I shows performance in execution time and energy consumption obtained by executing the algorithm on different configurations of the target architecture. A schematic block diagram of the algorithm is shown in Fig. 3. 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

TABLE I

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	cyk(k) ^a	E(μ J) ^c	cyk(k) ^a	su ^b	E(μ J) ^c	cyk(k) ^a	su ^b	E(μ J) ^c
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

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

TABLE II

CURRENT CONSUMPTION OF THE BOARD COMPONENTS IN THE DIFFERENT OPERATIONAL STATES

Operating Mode	Processing on Mr. Wolf @ 1.8 V	Digital Section @ 1.8 V	Analog Section @ 2.7 V	Battery Drain @ 3.7 V
Sleep	55 μ A	10 μ A	10 μ A	50 μ A
Streaming	55 μ A	7.2 mA	2.4 mA	6.4 mA
Application	1.0 mA	0.7 mA	2.4 mA	3.2 mA

overall performance (17.9 \times) because of the dominance of the MAP+ENCS kernel.

C. Power Consumption

To evaluate the performance of the architecture we set the operating frequency of Mr. Wolf to its most efficient operative point, 100 MHz at 0.8 V.

Table I 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 μ 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 μ 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 μ 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 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 are acquired at a sampling frequency of 1 KHz, and a new window of data 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, we power off the cluster and part of the SoC (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 II, 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 [21], [22], [13], 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.

IV. CONCLUSION AND FUTURE WORK

In this paper, we presented a complete system for wearable sensing and processing of biosignals, suitable for HMI design based on hand gesture recognition. 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, 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, outperforming the State-of-the-Art systems which reach only 11 h of operation with a standard 100 mAh LiPo battery. This demonstrates the capabilities of BioWolf, throwing the pillars for the next generation of unobtrusive and real-time embedded architecture for biosignal processing.

ACKNOWLEDGMENT

This work has been partially supported by the European H2020 FET project OPRECOMP under Grant 732631

REFERENCES

- [1] R. Meattini, S. Benatti, U. Scarcia, D. De Gregorio, L. Benini, and C. Melchiorri, "An semg-based human-robot interface for robotic hands using machine learning and synergies," *IEEE Transactions on Components, Packaging and Manufacturing Technology*, 2018.
- [2] S. Benatti, B. Milosevic, M. Tomasini, E. Farella, P. Schoenle, P. Bunjak, G. Rovere, S. Fateh, Q. Huang, and L. Benini, "Multiple biopotentials acquisition system for wearable applications," in *BIODEVICES*, 2015, pp. 260–268.
- [3] T. Starner, J. Weaver, and A. Pentland, "Real-time american sign language recognition using desk and wearable computer based video," *IEEE Transactions on pattern analysis and machine intelligence*, vol. 20, no. 12, pp. 1371–1375, 1998.
- [4] T. S. Saponas, D. S. Tan, D. Morris, R. Balakrishnan, J. Turner, and J. A. Landay, "Enabling always-available input with muscle-computer interfaces," in *Proceedings of the 22nd annual ACM symposium on User interface software and technology*. ACM, 2009, pp. 167–176.
- [5] Ottobock, <https://www.ottobock.com/prosthetics/upper-limb-prosthetics/solution-overview/myoelectric-prosthetics/>, 2018.
- [6] M. A. Oskoei, H. Hu *et al.*, "Support vector machine-based classification scheme for myoelectric control applied to upper limb," *IEEE Trans. Biomed. Engineering*, vol. 55, no. 8, pp. 1956–1965, 2008.

- [7] M. Atzori, M. Cognolato, and H. Müller, "Deep learning with convolutional neural networks applied to electromyography data: A resource for the classification of movements for prosthetic hands," *Frontiers in neurorobotics*, vol. 10, p. 9, 2016.
- [8] R. Meattini, S. Benatti, U. Scarcia, L. Benini, and C. Melchiorri, "Experimental evaluation of a semg-based human-robot interface for human-like grasping tasks," in *Robotics and Biomimetics (ROBIO), 2015 IEEE International Conference on*. IEEE, 2015, pp. 1030–1035.
- [9] B. Milosevic, S. Benatti, and E. Farella, "Design challenges for wearable emg applications," in *2017 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2017, pp. 1432–1437.
- [10] T. Labs, "Thalnic's myo armband," <https://www.myo.com/> year = 2013–2017.
- [11] J. Liu, F. Zhang, and H. H. 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*. IEEE, 2014, pp. 4095–4098.
- [12] X. Zhang, H. Huang, and Q. 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*. IEEE, 2013, pp. 5926–5929.
- [13] P. Gentile, M. Pessione, A. Suppa, A. Zampogna, and F. Irrera, "Embedded wearable integrating real-time processing of electromyography signals," in *Multidisciplinary Digital Publishing Institute Proceedings*, vol. 1, no. 4, 2017, p. 600.
- [14] M. Rossi, S. Benatti, E. Farella, and L. Benini, "Hybrid emg classifier based on hmm and svm for hand gesture recognition in prosthetics," in *Industrial Technology (ICIT), 2015 IEEE International Conference on*. IEEE, 2015, pp. 1700–1705.
- [15] A. Pullini, D. Rossi, I. Loi, A. D. 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)*, Sept 2018, pp. 274–277.
- [16] P. Kanerva, "Hyperdimensional computing: An introduction to computing in distributed representation with high-dimensional random vectors," *Cognitive Computation*, vol. 1, no. 2, pp. 139–159, 2009. [Online]. Available: <http://dx.doi.org/10.1007/s12559-009-9009-8>
- [17] P. D. 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*, Sept 2017, pp. 1–8.
- [18] A. Rahimi, S. Benatti, P. Kanerva, L. Benini, and J. M. Rabaey, "Hyperdimensional biosignal processing: A case study for EMG-based hand gesture recognition," in *IEEE International Conference on Rebooting Computing*, October 2016.
- [19] F. Montagna, A. Rahimi, S. Benatti, D. Rossi, and L. 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*, ser. DAC '18. New York, NY, USA: ACM, 2018, pp. 111:1–111:6. [Online]. Available: <http://doi.acm.org/10.1145/3195970.3196096>
- [20] M. R. Nuwer, G. Comi, R. Emerson, A. Fuglsang-Frederiksen, J.-M. Guérit, H. Hinrichs, A. Ikeda, F. J. C. Lucas, and P. Rappelsburger, "Ifcn standards for digital recording of clinical eeg," *Clinical Neurophysiology*, vol. 106, no. 3, pp. 259–261, 1998.
- [21] X. Liu, J. Sacks, M. Zhang, A. G. Richardson, T. H. Lucas, and J. 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*, vol. 64, pp. 1257–1261, 2017.
- [22] S. Benatti, G. Rovere, J. Bsser, F. Montagna, E. Farella, H. Glaser, P. Schuele, 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)*, June 2017, pp. 139–144.