



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

ARCHIVIO ISTITUZIONALE DELLA RICERCA

Alma Mater Studiorum Università di Bologna Archivio istituzionale della ricerca

Sub-100 μ W Multispectral Riemannian Classification for EEG-Based Brain-Machine Interfaces

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

Published Version:

Wang, X., Cavigelli, L., Schneider, T., Benini, L. (2021). Sub-100 μ W Multispectral Riemannian Classification for EEG-Based Brain-Machine Interfaces. IEEE TRANSACTIONS ON BIOMEDICAL CIRCUITS AND SYSTEMS, 15(6), 1149-1160 [10.1109/tbcas.2021.3137290].

Availability:

This version is available at: <https://hdl.handle.net/11585/963485> since: 2024-02-28

Published:

DOI: <http://doi.org/10.1109/tbcas.2021.3137290>

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)

Sub-100 μ W Multispectral Riemannian Classification for EEG-based Brain–Machine Interfaces

Xiaying Wang, *Student Member, IEEE*, Lukas Cavigelli, *Member, IEEE*, Tibor Schneider, *Student Member, IEEE*, Luca Benini, *Fellow, IEEE*

Abstract—motor imagery (MI) brain–machine interfaces (BMIs) enable us to control machines by merely thinking of performing a motor action. Practical use cases require a wearable solution where the classification of the brain signals is done locally near the sensor using machine learning models embedded on energy-efficient microcontroller units (MCUs), for assured privacy, user comfort, and long-term usage. In this work, we provide practical insights on the accuracy-cost trade-off for embedded BMI solutions. Our multispectral Riemannian classifier reaches 75.1% accuracy on a 4-class MI task. The accuracy is further improved by tuning different types of classifiers to each subject, achieving 76.4%. We further scale down the model by quantizing it to mixed-precision representations with a minimal accuracy loss of 1% and 1.4%, respectively, which is still up to 4.1% more accurate than the state-of-the-art embedded convolutional neural network. We implement the model on a low-power MCU within an energy budget of merely 198 μ J and taking only 16.9 ms per classification. Classifying samples continuously, overlapping the 3.5 s samples by 50% to avoid missing user inputs allows for operation at just 85 μ W. Compared to related works in embedded MI-BMIs, our solution sets the new state-of-the-art in terms of accuracy-energy trade-off for near-sensor classification.

Index Terms—brain–machine interface, motor imagery, edge computing, parallel computing, machine learning.

I. INTRODUCTION

BRAIN–machine interfaces (BMIs) provide a direct communication pathway between the human brain and an external device, such as a wheelchair [1], a prosthetic arm [2], or a drone [3]. They are especially useful for individuals with physical disabilities to regain independence [1], [4], and are often used in biomedical and clinical scenarios such as stroke rehabilitation [5]. Particularly interesting is the motor imagery (MI) BMI, where the subject’s intention is decoded from the brain activities recorded when the movement of a body part, e.g., feet, is merely imagined. The most common non-invasive technology to acquire brain signals is based on electroencephalography (EEG) thanks to its portability and relatively inexpensive hardware setup. A growing number of commercial EEG devices have been recently announced,

making BMIs increasingly common also outside laboratory and clinical environments [6], [7].

With the rising popularity of wearable BMI devices, it is crucial for a safe and successful BMI to meet key requirements, such as accurate classification, user comfort, long-term usage, real-time response, and, last but not least, privacy preservation. Traditional BMI systems collect and send the user data to a connected computer, a cloud, or a gateway with a wired or wireless setup, and often adopt offline, remote data processing to extract useful information. Connected cables yield a bulky BMI setup limiting subjects’ degrees of movement freedom. While wireless technology allows better user comfort, it introduces major drawbacks such as short battery life and privacy concerns [8]. Moreover, offline analysis leads to long delays in the information extraction. Recently, academic and industrial researchers have shared increasing interest in the concept of edge computing where the data is not remotely transmitted but directly processed near the sensor node [9], [10]. The collected sensitive biomedical signals are directly processed on-the-fly near the sensor right where the data is acquired and only the outcome of the processing is delivered, effectively curtailing communication-induced latency and eliminating most security and privacy risks introduced by the transmission and a remote processing pipeline.

However, many challenges need to be addressed in edge processing for BMI devices. Analyzing a large amount of sensor data at the edge requires miniaturized processing engines that provide enough compute power to extract useful information in real-time, while at the same time consume low power and are energy-efficient for a prolonged battery life [8]. A popular family of microcontroller units (MCUs) is the ARM Cortex-M series. They are commercialized in a wide variety ranging from the small and very-low power Cortex-M0 to the fast and high-performance Cortex-M7 devices. Another recently introduced edge-processing approach is the parallel ultra-low power (PULP) platform based on the open-source RISC-V instruction set architecture (ISA) [11]. Previous studies have proven that PULP processors outperform the Cortex-M series in terms of low power consumption, energy efficiency, and high compute capabilities [12], also for MI-BMI applications [13]. Moreover, the Mr. Wolf processor [14] of the PULP family has been deployed on a miniaturized BMI system, called Biowolf [15]. Thanks to the eight low-power, parallel processing units of Mr. Wolf, the Biowolf platform is one of the leading-edge systems for BMI applications. A next-generation PULP processor has been very recently released, called Vega [16]. Compared to Mr. Wolf, it provides more

Manuscript received August 31, 2021; revised October 30, 2021; accepted December 08, 2020.

X. Wang and L. Benini are with the Integrated Systems Laboratory, Department of Information Technology and Electrical Engineering (D-ITET), ETH Zürich, Switzerland (e-mails: {xiaywang, benini}@iis.ee.ethz.ch). L. Cavigelli is with Huawei Technologies, Computing Systems Lab, Zurich Research Center, Switzerland (e-mail: lukas.cavigelli@huawei.com). T. Schneider worked previously at the Integrated Systems Laboratory and is currently with the Networked Systems Group, D-ITET, ETH Zürich, Switzerland (e-mail: scitbor@ethz.ch).

This project was supported by the Swiss Data Science Center (SDSC) Ph.D. Fellowship under grant ID P18-04.

performance and better energy efficiency with its nine parallel cores, four floating point units (FPUs), and more on-chip memory. No work has yet demonstrated its capabilities with a real application.

Resource-efficient yet accurate classification algorithms are also crucial for a successful smart wearable device [8], especially for EEG applications where the high variability across subjects and among different recording sessions poses big challenges [9]. Deep neural networks (DNNs) have demonstrated impressive results in many fields, from outperforming humans in computer vision [17], vastly improving solutions for image processing [18] and speech recognition [19], natural language processing [20], large-scale recommender systems [21], and data analysis for sensors such as radar [22] that are not directly understandable by humans. Especially convolutional neural networks (CNNs) are largely applied in the image domain, but also for MI-BMI classification achieving state-of-the-art (SoA) accuracy [23]–[26]. However, they tend to grow in numbers of parameters making them often unfit for deployment on low-power low-cost MCUs, and require a large amount of training data to prevent overfitting. Whereas for BMI applications the acquisition and the labeling of EEG data are expensive, time-consuming, and prone to errors, resulting in scarce amounts of data available for training complex models with large numbers of parameters [27]. Over the years, successful methods have been proposed to extract discriminative, domain-specific features from EEG signals. The well-known common spatial patterns (CSP) algorithm learns spatial filters that discern between different MI tasks [28]. An improved algorithm, called filter-bank common spatial patterns (FBCSP), that accounts for multiple frequency bands has achieved better accuracy [29], [30]. More recent studies have proposed Riemannian methods to extract more comprehensive features also in absence of labeled data [31], [32]. The unsupervised feature calibration enables online adaptation of the classifier to combat the large inter-session variance in MI-BMIs [33]. So far, these methods are believed to be the most promising feature extractors for several kinds of BMI paradigms [34]–[36].

Few studies can be found in literature that have deployed MI-BMI models on edge devices. Bewalfi et al. [37] and Malekmohammadi et al. [38] have proposed CSP solutions based on field programmable gate arrays (FPGAs). While Schneider et al. [13] have deployed a compact CNN on Mr. Wolf achieving an accuracy of 71% on the popular BCI Competition IV-2a dataset and an energy consumption of 0.34 mJ, making it the embedded BMI with the lowest energy utilization. Wang et al. [39], on the other hand, have proposed an implementation of multispectral Riemannian classifier (MRC) on the same platform achieving 3% better accuracy while consuming 1.3 mJ, showing the trade-off between accuracy and cost.

In this work, we propose and release open-source¹ an improved implementation of MRC in terms of both accuracy and cost on the novel parallel platform Vega [16]. The main contributions are summarized as follows:




Fig. 1. Accuracy vs. memory footprint and computational complexity in terms of MAC operations of related work as well as the trade-off by our method (blue line) and the configuration used for most of our analyses (green). The memory footprint is estimated as the sum of the model size and the biggest consecutive feature maps based on the common layer-by-layer computation paradigm. As our method is not dominated by MAC operations as opposed to the competing DNN-based methods, we discuss execution time separately.

- We achieve an accuracy of 75.1% on a 4-class MI task when using only linear support vector machines (SVMs). The performance is further improved to 76.4% by tailoring the classifier to the subject within a given resource constraint, i.e., by tuning the hyperparameters of different types of classifiers, namely SVMs and multi-layer perceptrons (MLPs), for each subject.
- We quantize the feature extraction and the classifier of the MRC from full 32-bit float precision to a combination of 8-, 16-, 32-bit fixed- and floating-point representations, to maximize the efficiency on the hardware, while at the same time preserving similar accuracy. The quantization yields a 1.4% accuracy loss which is still 0.9% more accurate than the previous implementation of MRC [39] and 4.1% more accurate than the SoA embedded CNN with the lowest energy consumption [13].
- We implement for the first time a real application on the new Vega platform by demonstrating its potentials for low-power near-sensor analytics. We measure the runtime and power consumption on-board and find the configurations with the highest energy efficiency and the shortest execution time. Experimental measurements show that the former yields 16.9 ms and only 198 μ J and the latter takes only 7.9 ms and consumes 338 μ J per inference. Together with the high classification accuracy, it becomes the new SoA embedded BMI with the highest classification accuracy and the lowest energy utilization.

II. RELATED WORK

The next generation of smart edge BMIs have to manage the three-way trade-off among (a) algorithmic performance in terms of high classification accuracy, (b) cost in terms of low computational runtime and memory footprint, and (c) power

¹<https://github.com/pulp-platform/multispectral-riemannian>

TABLE I
COMPARISON TO RELATED WORKS ON BCI COMPETITION IV-2A DATASET
IN TERMS OF ACCURACY, MODEL SIZE (TOP), QUANTIZED ACCURACY,
COMPUTATION RUNTIME AND ENERGY PER INFERENCE (BOTTOM).

	Acc. [%]		Num. param.	
	2-	4-class		
S. ConvNet '17 [23]**	79.90	74.31	47.3 k	
EEGNet '18 [13], [24]	-	71.30	2.63 k	
FBCSP '18 [40]	-	73.7	261 k	
MRC '18 [40]	-	75.47	32.6 k	
MSFBCNN '19 [25]	-	75.8	155 k	
EEG-TCNet '20 [26] [†]	-	77.35	4.27 k	
MRC-MLP '20 [41]	-	76.1	5.7 M	
CSP+Riemannian '20 [42] [◊]	-	77.82	-	
JSTFD+LDA '20 [43]	79.6	-	-	
FB+3D-CNN '21 [44]	86.96	-	46 M	
TS-SEFFNet '21 [45]	-	74.71	282 k	
MRC-				
-Mr.Wolf '21 [39]	-	75.1	4.55 k	
-Vega[‡] (ours)	-	75.1	4.55 k	
	Quant. acc. [%]	Time [ms]	En. [μJ]	
WOLA-CSP '18 [37]	78.85	-	430	301000
CSP+SVM '19 [38]	80.55	67.21	11.66	978
Q-EEGNet '20 (50 [♯]) [13]	-	70.90	28.67	337
Q-EEGNet '20 (350 [♯]) [13]	-	70.90	13.64	678
MRC-				
-Mr.Wolf '21 (100 [♯]) [39]	-	74.1	33.39	1304
-Vega (350[♯]) (ours)	-	74.1	7.9	338
-Vega (50/160[‡]) (ours)	-	74.1	16.9	198

** Respectively reproduced in [44] and [26] for 2- and 4-class tasks.

[†] Improves to 83.84% with subject-specific variable network architectures and hyperparameters.

[◊] With subject-specific feature selection based on EEG channels.

[‡] Improves to 76.4% with subject-specific variable classifier.

[♯] Clock frequencies in MHz. All cores run at the same frequency.

[‡] The FC core runs at 50MHz. The cluster cores run at 160MHz.

and energy consumption [8]. Many research works have been conducted to boost the performance of BMI algorithms in terms of classification accuracy, while only recently increasing attention has been paid to the resources that are required by these algorithms (see Table I). Considering the popular 4-class BCI Competition IV-2a dataset [46], the reproducible SoA accuracy is around 77% reached by EEG-TCNet [26], when no additional algorithmic optimizations are used, such as subject-specific feature extraction by tuning the model architecture or subject-specific channel selection. Compared to other existing works on the same dataset [23], [25], [40], [44], [45], EEG-TCNet requires orders of magnitude fewer resources in terms of memory footprint and computational complexity, while reaching higher classification accuracy, as depicted in Fig. 1. However, if we consider the common layer-by-layer computation paradigm for the classification inference, the required memory footprint is out-of-reach for the fast on-chip memory of ultra-low power MCUs such as Mr. Wolf with 512 kB L2 and 64 kB L1 memory.

Bewalfi et al. [37] are the first ones to propose an embedded MI-BMI based on CSP algorithms, called WOLA-CSP, on a Stratix-IV FPGA. The achieved accuracy is 78.85% on a 2-class task, while the average power consumption is 0.7 W

with an execution runtime of 430 ms, yielding an energy consumption of 301 mJ. Malekmohammadi [38] have proposed a much more energy efficient solution based on CSP and SVM on a Virtex-6 FPGA. The 2-class and 4-class classification accuracy scores are respectively 80.55% and 67.21%, with an energy consumption of 0.978 mJ. EEGNet [24], on the other hand, is a very compact CNN with only a few thousand model parameters. It has been successfully quantized with Q-EEGNet [13] and implemented on Mr. Wolf reaching an energy consumption of only 0.678 mJ. It has proven to be three orders of magnitude more energy efficient than a EEGNet implementation on commercially available MCUs based on ARM Cortex-M architecture [47], making it the SoA embedded CNN in terms of energy efficiency and compact model size, as reported in Table I. However, the classification performance is around 6% less accurate than the EEG-TCNet [26]. Another very promising algorithm to extract discriminant features from MI EEG data is based on Riemannian geometry [35], [41]. Hersche et al. [40] have proposed a multiscale temporal and spectral Riemannian classifier with SVM and achieved a 4-class accuracy of 75.1% outperforming both EEGNet and CSP-based models by around 3%–7%. It has been further optimized in terms of resource usage by using only one temporal window, quantized to mixed-precision representations, and efficiently implemented on Mr. Wolf with an energy consumption of 1.304 mJ and a runtime of 33.39 ms [39]. It is the most accurate embedded solution, however, it comes at the cost of a higher energy consumption and a longer runtime compared to [13], [38].

In this work, we further improve the classification accuracy of MRC and implement it on a more energy-efficient MCU called Vega. Similar to Mr. Wolf, Vega is a parallel multi-core processor based on the open-source RISC-V ISA with custom extensions, but it has an improved hardware architecture and more resources, yielding 3.23× and 4.39× better efficiency compared to Mr. Wolf, respectively for integer and floating-point operations [16].

III. MODEL DESIGN

The MRC proposed in this work is based on [39]. Non-linear features are first extracted from different frequency bands using the Riemannian covariance method [35]. Subsequently, a machine learning classifier is trained to classify different classes.

A. Feature Extraction

Fig. 2 illustrates the proposed model. The input EEG data \mathbf{X} has dimensions number of EEG channels $N_{ch} = 22$ times number of time samples $N_s = 875$ and is filtered with $n = 18$ infinite impulse response (IIR) bandpass filters to obtain the filtered signals $\hat{\mathbf{X}}_k$:

$$\hat{\mathbf{X}}_k = \text{IIR } b_k(\mathbf{X}), \quad (1)$$

with $k \in \{1, 2, \dots, n\}$. They are realized as 4th order Butterworth filters with the frequency bands starting from 4 Hz until 40 Hz with a bandwidth of 2 Hz each, i.e, the lower and upper cut-off frequencies of the IIR filters are 4 Hz and 6 Hz,




Fig. 2. multispectral Riemannian classifier with $n = 18$ frequency bands and one time window. The regions surrounded by dashed lines represent the three parts of the computation. The shaded areas represent regions which are processed by all 9 cores in parallel. Adapted from [39].




Fig. 3. Mixed-precision, quantized MRC of a single frequency band, showing the representation of each intermediate signal [39].

6 Hz and 8 Hz, 8 Hz and 10 Hz, ..., until 38 Hz and 40 Hz. It has been proven that these frequency bands yield the best performance [39], [40].

Afterwards, the spatial covariance matrix $\mathbf{C}_k \in \mathbb{R}^{N_{ch} \times N_{ch}}$ for each frequency band can be estimated as

$$\mathbf{C}_k = \hat{\mathbf{X}}_k \hat{\mathbf{X}}_k^T + \rho \mathbf{I}, \quad (2)$$

with ρ being the regularization parameter that scales the identity matrix \mathbf{I} . These covariance matrices are by construction symmetric and, with sufficient data, positive definite. A correct way to manipulate them for the final classifier is to rely on Riemannian geometry [48]. According to which, each covariance matrix can be projected locally onto an Euclidean tangent plane, which is done by the next operation:

$$\mathbf{W}_k = \mathbf{C}_{ref,k}^{-1/2} \mathbf{C}_k \mathbf{C}_{ref,k}^{-1/2}. \quad (3)$$

The covariance matrix is transformed by multiplying it with a reference matrix $\mathbf{C}_{ref,k}^{-1/2}$, which is obtained beforehand during the model training by averaging the covariance matrices of all the training data for each frequency band. The mapping of the covariance matrix to the tangent space can be interpreted as a whitening operation [48]; hence, we name it *Whitening* in this paper as a shorthand.

The matrix logarithm is subsequently applied on the *whitened* matrices \mathbf{W}_k for finalizing the data mapping:

$$\mathbf{L}_k = \text{logm}(\mathbf{W}_k). \quad (4)$$

The function $\text{vect}(\mathbf{L}_k)$ follows and vectorizes the output \mathbf{L}_k of the matrix logarithm. \mathbf{L}_k is a symmetric matrix and its upper right off-diagonal elements are multiplied by $\sqrt{2}$ to preserve the norm. The diagonal values $l_{k,(i,i)}$ and the upper-right off-diagonal elements $l_{k,(i,j)}$, with $i \in 1, 2, \dots, N_{ch}$ and $j \in 2, 3, \dots, N_{ch}$, are concatenated to form a vector:

$$\mathbf{v}_k = \text{vect}(\mathbf{L}_k) = [l_{k,(i,i)}] \oplus [\sqrt{2} \cdot l_{k,(i,j)}]. \quad (5)$$

The features vectors \mathbf{v}_k from all frequency bands are finally concatenated and fed as input to the classifier.

B. Classification

Based on the obtained feature vector, a simple classifier then decides whether the person was thinking about moving the ‘left hand’, ‘right hand’, ‘both feet’, or ‘tongue’, using the BCI

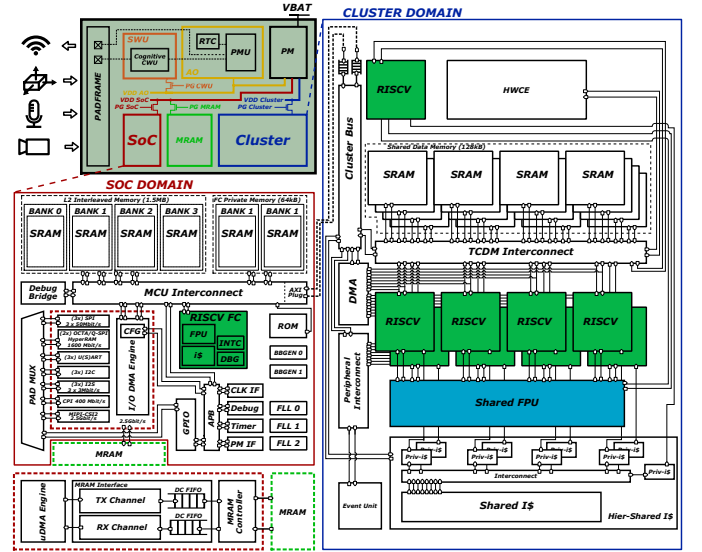


Fig. 4. SoC architecture and power domains of Vega. Adapted from [16].

Competition IV-2a dataset [46]. We analyze two scenarios, a one-fits-all solution and a trade-off analysis of the average accuracy against the memory requirements.

For the former, we train a linear SVM on the features extracted for each patient’s training data and evaluate the accuracy of the overall pipeline with and without quantization (following the quantization described in Sec. V).

For the latter, we search for linear SVMs and MLPs for each patient, tuning either the SVM’s L2 regularization parameter or the MLP’s number of hidden layers (between 1 and 3), the activation function (ReLU or sigmoid), and the number of nodes for each hidden layer separately in $\{2, 4, 8, \dots, 2048\}$.

IV. PARALLEL ULTRA-LOW POWER PROCESSOR

In this section we summarize the key features of the PULP MCU used in this work, called Vega [16]. It is designed for near-sensor applications featuring two main power domains, i.e., SoC and cluster, with 10 RISC-V cores, in addition to an always-on domain with a *cognitive* wake-up unit, as depicted in Fig. 4. It operates from 0.6 V – 0.8 V and can be scaled from a fully retentive *cognitive* sleep mode consuming only 1.7 μW up to 32.2 GOPS peak performance consuming 49.4 mW.

Vega features one RISC-V core in the SoC domain, called fabric controller (FC), and nine parallel cores in the cluster. The FC handles the several peripherals of the MCU and executes simple computations. For compute-intensive tasks, the cluster is activated and the tasks are offloaded to the nine parallel cores with four-pipeline stages. Four shared FPUs are available in the cluster for FP32, FP16, and bfloat floating-point operations, while a multi-banked L1 memory of size 128 kB is shared among the cores and connected through a 1-cycle tightly coupled data memory (TCDM) interconnect for fast access. A hierarchical program cache serves the cores with 512 B private per-core plus 4 kB of 2-cycle-latency shared cache to maximize efficiency. A direct memory access (DMA) unit, orchestrated by one of the nine cores, can autonomously transfer data between the L1 memory and the much larger L2 memory present in the SoC, which is a multi-banked static random-access memory (SRAM) of size 1.5 MB and 64 kB private memory for the FC. An additional DMA engine in the SoC can independently transfer I/O data between the L2 in the SoC and the 4MB of non-volatile magnetoresistive random-access memory (MRAM) and an eventual off-chip HyperRAM. Thanks to the DMA units, the computation and the data transfer can be fully overlapped.

The RISC-V cores are based on RVC32IMF ISA with additional custom extensions named Xpulp for near-sensor analytics algorithms. The extensions include hardware loops, post-incremental load and store, 2-way and 4-way single instruction, multiple data (SIMD) operations respectively for 16-bit and 8-bit data types. Together with the 9-core parallelism and optimized memory hierarchy, Vega can deliver up to 15.6 8-bit GOPS and up to 612 GOPS/W, and up to 3.3 GFLOPS and 129 GFLOPS/W [16].

V. QUANTIZATION AND IMPLEMENTATION

Unlike the feed-forward CNNs, which can be quantized to very low numerical precision, e.g., down to 8 bits or beyond without accuracy loss [49]–[51], the quantization of the MRC may cause significant accuracy degradation, mainly due to the numerical stability of the IIR filters and the computation of the matrix logarithm. The accuracy has to be preserved, while at the same time lower precision yields less memory footprint and more energy-efficient implementation by exploiting the 2- and 4-way SIMD instructions of the hardware architecture. Hence, a mixed-precision quantization is proposed with a combination of 8-, 16-, and 32-bit fixed-point and 32-bit floating-point representations, as depicted in Fig. 3. Moreover, to fully exploit the multi-core nature of the Vega MCU, we implement the MRC by splitting the feature extraction into two separate blocks. As Fig. 2 highlights, the features of every frequency band can be computed individually. We divide the computation between the 9 cores as follows:

- 1) The first part of the feature extraction (in blue) computes an IIR filter, followed by several matrix operations. Each IIR filter cannot be parallelized, as its implementation keeps state while moving along the signal. Instead, we can perform the filtering of the N_{ch} different EEG channels in parallel. The subsequent matrix operations

can then be computed using their optimized, parallel implementations. Hence, we use all 9 cores to compute the whitened signal W_k for a single frequency band k , and repeat this for all $n = 18$ bands.

- 2) The second part of the feature extraction (in red) requires computation of the Eigendecomposition (EVD). Parallel implementations of such algorithms do not yield a speedup factor close to the theoretical maximum. Hence, we utilize the multi-core nature of the MCU by computing multiple decompositions in parallel, each of them assigned to a single processing core.
- 3) The final classification (in green) is computed on a single core since it accounts for a negligible part of the execution time.

In the following paragraphs, we elaborate on the details of the quantization and the implementation on Vega.

A. Quantization Terminology

Assume a value $x \in \mathbb{R}$, the fixed-point representation $\tilde{x} \in \mathbb{Z}$ (in two's complement) of x is written as

$$\tilde{x} = \text{clip} \left(\text{round} \left(x \cdot \frac{R_x}{S_x} \right), -R_x, R_x - 1 \right) \quad (6)$$

with $R_x = 2^{n_x - 1}$, or simplified, ignoring the clip and round function,

$$\tilde{x} = x \cdot \frac{R_x}{S_x}. \quad (7)$$

S_x represents the *dynamic range* of x which can be represented by \tilde{x} . If $|x| > S_x$, then \tilde{x} is clipped. Choosing the dynamic range S_x to be a power of two results in a traditional fixed-point format. R_x represents the number of quantization levels. If \tilde{x} is represented with n_x bits, then $R_x = 2^{n_x - 1}$. The fraction R_x/S_x is called the *conversion factor*.

B. IIR Bandpass Filters

The input data \mathbf{X} is quantized to 8 bits and is filtered with IIR filters. Unlike finite impulse response (FIR) filters, i.e., convolutions, that are always numerically and asymptotically stable, the IIR filters can become unstable, especially when they are quantized. The internal accumulators can diverge, even if the output remains bounded. Different realization forms exist [52], yielding different asymptotic behavior of the internal signals and thus influencing the numerical stability. We implement the Direct-Form I, shown in Fig. 5. It ensures that no internal overflows happen, because all internal registers store either the input or the output of the filter [52]. A typical approach for quantizing an IIR filter is to express them as a cascade of second-order sections (SOSs). In this representation, multiple separate second-order IIR filters are cascaded allowing different sections to be quantized with different dynamic ranges, thus minimizing the effect of quantizing the filter coefficients on the impulse response. We implement the 4th order Butterworth filters as a cascade of two SOSs with the following transfer function:

$$H(z) = \prod_{m=0}^{M-1} \frac{b_{0,m} + b_{1,m}z^{-1} + b_{2,m}z^{-2}}{1 + a_{1,m}z^{-1} + a_{2,m}z^{-2}}, \quad (8)$$




Fig. 5. IIR filter realization: Direct Form I as defined in [52].

where z is a complex variable, M is the number of sections, a and b are respectively the reverse and forward coefficients as depicted in Fig. 5. There are a total of $18 \cdot 2 \cdot 5 = 180$ coefficients, which we release together with the code¹.

The choice of the quantization levels depends on three requirements: a) The quantization must not change the frequency response of the filters; b) No numerical overflow must happen; c) The underlying hardware can be optimally used. Quantizing all filter coefficients with 8 bits has a significant impact on the impulse response of the filter, as shown in Fig. 7 and Fig. 8, where we observe overshooting, undershooting, or shifting behaviours of the frequency responses compared to the full-precision ones. We increase the number of bits and observe that a 12-bit representation has unnoticeable effect on the frequency responses which closely overlap with the full-precision ones. Another strictly related parameter is the number of bits used for the internal signals. We observe that using 12 bits for the filter coefficients and 16-bit registers to accumulate intermediate values does not cause any overflow for the dataset of interest, and it allows the usage of 2-way SIMD operations for the following iteration of the filter computation yielding a more efficient implementation. Thus, we choose 12 bits to quantize the filter coefficients, since a higher number of bits is not necessary for preserving the frequency responses and use 16 bits to represent the intermediate results of the SOS, as depicted in Fig. 6. Each SOS contains three multiply-and-accumulates (MACs) for the forward accumulation and two MACs for the backward accumulation. This enables the usage of SIMD instructions with bit-width 16.

Moreover, the dynamic ranges of the coefficients of each SOS are chosen independently to minimize the quantization error and set to be a power of two, such that no expensive divisions are required with the advantage of using more energy-efficient bit-shift operations. The RISCY cores on Vega support instructions (namely `p.addNR` and `p.subNR`), which normalize and round the result of an addition or a subtraction accordingly. Therefore, all intermediate results can be rounded efficiently, improving the precision of the filter. For the parallel computation, we assign each 1D signal of size N_s from the N_{ch} EEG electrodes to individual cores for the filtering. In other words, $\tilde{\mathbf{X}}_k$ is computed in parallel using the nine cores. Once finished, $\tilde{\mathbf{X}}_{k+1}$ is subsequently assigned to the cluster for the parallel computation.

C. Covariance Matrix




Fig. 6. IIR filter realization and implementation. Diagram showing the registers and the operations of a single SOS, where arrows in blue represent register transfers and shuffle operations, and red arrows represent MAC instructions. For an IIR filter with two SOS the output $\tilde{x}^{(1)}$ can be reused.

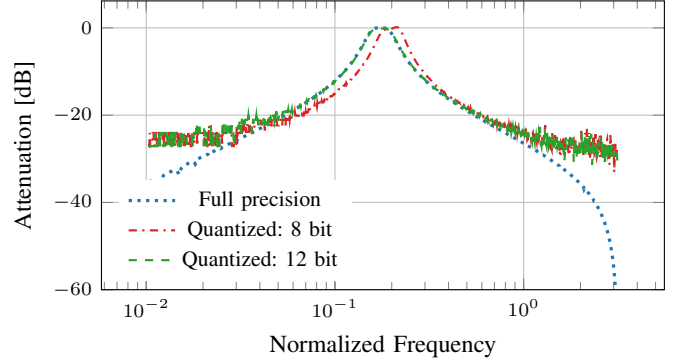


Fig. 7. Frequency response of a 4th order Butterworth IIR band-pass filter with lower and upper cut-off frequencies of 6Hz and 8Hz, respectively. The full precision filter (blue dotted line) is compared with the quantized implementations, where the filter coefficients are quantized to 8 (red dash-dotted line) and 12 (green dashed line) bits, and the output to 8 bits.

Following the terminology introduced in Sec. V-A, the covariance matrix \mathbf{C}_k is quantized as

$$\tilde{\mathbf{C}}_k = \frac{\tilde{\mathbf{X}}_k^T \tilde{\mathbf{X}}_k + \tilde{\rho} \mathbf{I}}{\frac{R_c^2 S_c}{S_x^2 R_c}} = \frac{\tilde{\mathbf{X}}_k^T \tilde{\mathbf{X}}_k + \tilde{\rho} \mathbf{I}}{F_c}, \quad (9)$$

with $F_c = \frac{2^{2n_x - 2} \cdot S_c}{2^{n_c - 1} \cdot S_x^2}$, $\tilde{\mathbf{X}}_k$ being the filtered signals represented in 8 bits, R the quantization levels, and S the dynamic range. We force S_c/S_x^2 to be a power of two, such that the scaling factor F_c is also a power of two, allowing us to implement the transformation with a bit-shift. Note that $\tilde{\mathbf{C}}_k$ is a symmetric matrix, hence we compute only the upper-right triangle and copy the remaining elements. 4-way SIMD operations are used for the 8-bit numbers and the computation is implemented concurrently by splitting the upper-right part of the output matrix among all processing units.

D. Whitening

We notice that the input to the matrix logarithm is sensitive to quantization errors due to the calculation of the EVD [53]. Thus, to avoid accuracy degradation, we use the high dynamic range available with 32 bits for \mathbf{W}_k . Let n_c and n_{ref} be the number of bits to represent $\tilde{\mathbf{C}}_k$ and $\tilde{\mathbf{C}}_{\text{ref},k}^{-1/2}$, respectively. We rescale $\tilde{\mathbf{C}}_k$ to use the entire dynamic range with n_c bits, while the reference matrix $\mathbf{C}_{\text{ref},k}^{-1/2}$ is quantized to n_{ref} bits during training. We do not scale either the intermediate, nor the final result of Eq. 3 to minimize the quantization error. Since we can exploit either 4- or 2-way SIMD operations, we test both $n_c = n_{\text{ref}} = 8$ and 16. The accuracy drops significantly with the former, while the latter causes overflows. Hence, we reduce




Fig. 8. A zoomed-in illustration of the frequency responses. The quantization to 8 bits causes overshoot (blue), undershoot (orange), or shift (red and green) of the frequency bands. The quantization to 12 bits has minimal effect compared to the full precision filters. Four out of 18 filters are illustrated, while the others present similar behavior.

n_{ref} , until training completes without overflow, resulting in $n_{\text{ref}} = 11$. With $n_c = 16$ and $n_{\text{ref}} = 11$ we obtain similar accuracy to the full-precision version. The dynamic range S_w is determined by

$$\tilde{\mathbf{W}} \frac{R_w}{S_w} = \tilde{\mathbf{C}}_{\text{ref}}^{-1/2} \tilde{\mathbf{C}} \tilde{\mathbf{C}}_{\text{ref}}^{-1/2} \frac{R_c R_{\text{ref}}^2}{S_c S_{\text{ref}}^2}, \quad (10)$$

which means

$$S_w = \frac{2^{n_w-1} \cdot S_c S_{\text{ref}}^2}{2^{n_c-1} \cdot 2^{2n_{\text{ref}}-2}}. \quad (11)$$

The result of the Whitening block is then multiplied by $S_w/2^{n_w-1}$ to be converted to floating-point representation as input to the matrix logarithm.

We implement the two matrix-matrix multiplications (MMMs) of the Whitening operation using the PULP-DSP library [54]. The first multiplication is computed in 16 bits by exploiting the SIMD instructions, the second in 32 bits to preserve the full dynamic range. The DSP library splits the computation into small chunks of size 2×4 in the 16-bit implementation, and 2×2 in the 32-bit version. Each chunk is assigned to a single core to be computed.

E. Matrix Logarithm

Algorithm 1 Householder Tridiagonalization [55]

```

1: procedure HHTRIDIAG( $\mathbf{A}$ )  $\triangleright \mathbf{A} \in \mathbb{R}^{n \times n}$  with elements  $a_{j,k}$ 
2:    $\mathbf{Q} \leftarrow \mathbf{I}$ 
3:   for  $k \in \{1, 2, \dots, n-2\}$  do  $\triangleright$  working on column  $k$ 
4:      $\alpha \leftarrow -\text{sgn}(a_{k+1,k}) \cdot \|\mathbf{A}_{k+1:n,k}\|$   $\triangleright$  euclidean norm
5:      $r \leftarrow \sqrt{\frac{1}{2}(\alpha - a_{k+1,k})}$ 
6:      $\vec{v} \leftarrow [0, \dots, 0, a_{k+1,k} - \alpha, a_{k+2,k}, \dots, a_{n,k}]^T / 2r$ 
7:      $\mathbf{P} \leftarrow \mathbf{I} - 2\vec{v}\vec{v}^T$   $\triangleright$  Householder matrix
8:      $\mathbf{A} \leftarrow \mathbf{P}\mathbf{A}\mathbf{P}^T$ 
9:      $\mathbf{Q} \leftarrow \mathbf{Q}\mathbf{P}$ 
10:  end for
11:  return  $\mathbf{A}, \mathbf{Q}$   $\triangleright \mathbf{A}$  tridiagonal,  $\mathbf{Q}$  orthogonal
12: end procedure

```

The matrix logarithm of a square matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ is defined in terms of its EVD, as

$$\text{logm}(\mathbf{A}) = \mathbf{Q}^{-1} \text{logm}(\mathbf{D}) \mathbf{Q}, \quad (12)$$

where $\mathbf{A} = \mathbf{Q}^{-1} \mathbf{D} \mathbf{Q}$, with \mathbf{D} containing the Eigenvalues of \mathbf{A} on its diagonal and \mathbf{Q} containing the Eigenvectors of \mathbf{A} . The logarithm of a diagonal matrix \mathbf{D} is computed by applying the logarithm to its diagonal element, i.e.,

$$\text{logm} \left(\begin{bmatrix} d_1 & 0 & \dots & 0 \\ 0 & d_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & d_n \end{bmatrix} \right) = \begin{bmatrix} \log(d_1) & 0 & \dots & 0 \\ 0 & \log(d_2) & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \log(d_n) \end{bmatrix}. \quad (13)$$

The matrix logarithm exists in $\mathbb{R}^{n \times n}$ only if $\mathbf{A} \succ 0$ is positive definite, because a positive definite matrix has only positive Eigenvalues which ensure the existence of logarithm in \mathbb{R} .

The feature extraction of the MRC only computes the matrix logarithm of dense and symmetric matrices, meaning that we can optimize the EVD by computing first the tridiagonal decomposition which produces a tridiagonal matrix similar to the original one, i.e., the Eigenvalues are preserved. Computing the EVD of a tridiagonal matrix yields less computational effort. The final transformation is of the form

$$\mathbf{A} = \mathbf{Q}_t^T \mathbf{T} \mathbf{Q}_t = \mathbf{Q}_t^T \mathbf{Q}_d^T \mathbf{D} \mathbf{Q}_d \mathbf{Q}_t, \quad (14)$$

where \mathbf{Q}_t is the orthogonal matrix for the tridiagonal transformation and \mathbf{Q}_d the one for the EVD. $\mathbf{Q}_d \mathbf{Q}_t$ is an orthogonal matrix containing the Eigenvectors of \mathbf{A} , and \mathbf{T} is of the form

$$\mathbf{T} = \begin{bmatrix} a_1 & b_2 & 0 & 0 & \dots & 0 & 0 \\ b_2 & a_2 & b_3 & 0 & \dots & 0 & 0 \\ 0 & b_3 & a_3 & b_4 & \dots & 0 & 0 \\ 0 & 0 & b_4 & a_4 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \dots & a_{n-1} & b_n \\ 0 & 0 & 0 & 0 & \dots & b_n & a_n \end{bmatrix} \quad (15)$$

containing the same Eigenvalues of \mathbf{A} .

To compute the tridiagonal matrix, we use the Householder transformation [55], reported in Algorithm 1. This algorithm has complexity $\mathcal{O}(n^4)$, since it contains MMMs (on line 8 and line 9), which take $\mathcal{O}(n^3)$, for a total of $\mathcal{O}(n)$ iterations. We can improve Algorithm 1 by grouping the operations differently:

$$\mathbf{A}^+ = \mathbf{P}\mathbf{A}\mathbf{P}^T = (\mathbf{I} - 2\vec{v}\vec{v}^T) \mathbf{A} (\mathbf{I} - 2\vec{v}\vec{v}^T) \quad (16)$$

$$= \mathbf{A} - 2\vec{v}\vec{v}^T \mathbf{A} - 2\mathbf{A}\vec{v}\vec{v}^T + 4\vec{v}\vec{v}^T \mathbf{A}\vec{v}\vec{v}^T \quad (17)$$

Recall that \mathbf{A} is symmetric. We can simplify this equation introducing the row vector $\vec{b}^T = \vec{v}^T \mathbf{A} \in \mathbb{R}^n$, the matrix $\mathbf{B} = \vec{v}\vec{b}^T \in \mathbb{R}^{n \times n}$ and the constant $c = \vec{b}^T \vec{v} \in \mathbb{R}$, and get:

$$\mathbf{P}\mathbf{A}\mathbf{P}^T = \mathbf{A} - 2(\mathbf{B} + \mathbf{B}^T) + 4c\vec{v}\vec{v}^T, \quad (18)$$

with $\vec{b}^T = \vec{v}^T \mathbf{A}$, $\mathbf{B} = \vec{v}\vec{b}^T$, $c = \vec{b}^T \vec{v}$. This shows, that we can reduce the two matrix multiplications ($\mathcal{O}(n^3)$) to one vector matrix multiplication, one inner vector product, two outer vector products, two scalar matrix multiplications, and three matrix additions, resulting in a total complexity of $\mathcal{O}(n^2)$. Similarly, we can compute the update for matrix \mathbf{Q} to

TABLE II
CLASSIFICATION ACCURACY (%) ON 2- (MARKED WITH **) AND 4-CLASS MI TASKS FOR EMBEDDED IMPLEMENTATION.

Ref.	WOLA-CSP	CSP+SVM		Q-EEGNet		MRC					
	[37]**	[38]**	[38]	[13]	[13]	[40]	[40]	ours	ours	ours	ours
Classifier	LDA	SVM	SVM	built-in	built-in	lin. SVM	lin. SVM	lin. SVM	lin. SVM	variable	variable
Precision	16-bit	20-bit	20-bit	full	8-bit	full	full	full	mixed	full	mixed
#features	–	–	–	200k	200k	32 637	10 879	4 554	4 554	4 554	4 554
#temp. win. t	–	–	–	–	–	3	1	1	1	1	1
#freq. bands f	–	–	–	–	–	43	43	18	18	18	18
Cov. Mat. Reg. ρ	–	–	–	–	–	0.0	0.0	1.0	1.0	1.0	1.0
Subject 1	86.8	91.7	79.2	81.0	81.0	90.0	91.8	91.8	90.7	svm: 91.8	svm: 90.8
Subject 2	63.9	59.7	56.3	57.6	53.1	55.5	51.6	53.7	51.2	mlp: 60.4	mlp: 57.6
Subject 3	94.4	95.8	87.5	87.9	91.2	81.3	83.5	83.5	81.0	svm: 83.5	svm: 81.0
Subject 4	68.8	77.1	63.9	61.6	58.1	71.9	73.3	73.7	74.1	mlp: 74.6	mlp: 74.1
Subject 5	56.3	67.4	41.0	70.6	68.4	69.6	63.4	68.8	63.0	svm: 68.8	mlp: 63.8
Subject 6	69.4	69.4	46.9	53.4	50.1	56.7	58.6	56.7	56.3	mlp: 57.2	svm: 56.3
Subject 7	78.5	78.5	81.9	75.7	75.2	85.6	86.7	84.1	85.9	svm: 87.0	svm: 85.9
Subject 8	97.9	97.2	76.0	77.4	81.2	83.8	81.6	81.5	82.7	mlp: 82.3	mlp: 83.4
Subject 9	93.8	88.2	72.2	76.7	79.7	84.9	82.6	82.2	81.8	mlp: 82.2	svm: 81.8
Avg. acc.	78.85	80.55	67.21	71.3	70.9	75.5	74.8	75.1	74.1	76.4	75.0
Std. dev.	–	–	–	11.5	14.3	12.8	13.9	12.2	13.2	11.3	12.0

TABLE III
COMPUTE TIME FOR MRC ON VEGA.

	baseline		improved EVD		parallel execution		parallel speedup	ops/cycle [†]
	[cycles]	(rel. %)	[cycles]	(rel. %)	[cycles]	(rel. %)		
Filter	5 637k	(15.73%)	5 637k	(33.96%)	779k	(33.41%)	7.24	4.45
Cov. matrix	3 071k	(8.57%)	3 071k	(18.50%)	393k	(16.86%)	7.81	10.14
Whitening	809k	(2.26%)	809k	(4.87%)	180k	(9.00%)	4.31	2.12
Matrix logm.	26 314k	(73.40%)	7 070k	(42.58%)	934k	(40.08%)	7.60	2.72
SVM	15k	(0.04%)	15k	(0.09%)	15k	(0.65%)	-	1.20
Total [cycles]	36 912k		17 365k		2 417k		7.18	
MACs/cycle [‡]	0.82		0.82		5.62			
FLOPs/cycle [◊]	0.10		0.36		2.72			
instructions/cycle	0.845		0.870		0.664			
Runtime [ms]*	233.5		115.7		16.9			
Avg. power [mW]*	5.3		5.2		11.7			
Energy [μ J]*	1238		602		198			

[‡] Number of fixed-point MACs over number of cycles w/o matrix logarithms.

[◊] FLOP/cycles during matrix logarithms, approximated by treating all floating-point instructions equally.

[†] MACs or FLOP per cycle for the concurrent implementation except SVM.

* With FC at 50 MHz, the cluster cores at 160 MHz, and 0.6 V supply. Cluster activation included (0.9 ms, 2 mW, 1.8 μ J).

MI classes from 9 different subjects. There are 288 trials for each of the training and testing sets. Each trial consists of a fixation cross on the screen, followed by a cue. Afterwards the subject starts the motor imagery corresponding to the cue. The data is sampled at 250 Hz and in this work we consider 3.5 s time window one second after the appearance of the cue as in [40], resulting in 875 time points per EEG channel, i.e., $N_s = 875$.

A. Accuracy, Memory Footprint, and Quantization

MRC can be scaled to use more or fewer frequency bands and temporal windows. Hersche et al. [40] have shown that $f = 43$ frequency bands and a single temporal window $t = 1$ can already achieve comparable accuracy (74.8% on average) to the full MRC (75.5%) while requiring $3\times$ fewer features.

In this work, we use only one temporal window $t = 1$ of 3.5 s and further scale down the number of frequency bands as in [39]. The results show that with $2.4\times$ less frequency bands, i.e. $f = 18$, of bandwidth 2 Hz between 4 and 40 Hz, the full precision model achieves slightly higher accuracy by introducing the regularization for the covariance matrix with the hyperparameter $\rho = 1$. We then gain an additional 1.3% in full precision by tuning the classifier and its hyperparameters to the subject, as shown in Table II. The achieved accuracy is lower than some of the SoA works reported in Table I, when considering only full precision models. Nevertheless, our MRC requires significantly less storage for the parameters than most of the related works at the cost of minor accuracy degradation, enabling the embedded implementation on low-power MCUs.

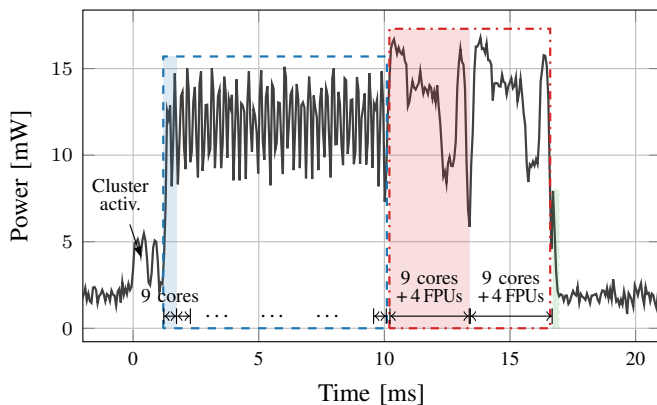


Fig. 10. Power measurement with the FC at 50 MHz and the cluster at 160 MHz. The colors match the compute blocks in Fig. 2 explained in Sec. V.

To efficiently deploy the proposed model, we proceed with quantization as described in Sec. V. This leads to a drop in accuracy of merely 1.0% from 75.1% to 74.1% when using a linear SVM for the classification step, or 1.4% from 76.4% to 75.0% when selecting the best classifier for each subject as described in Section III-B. The quantization allows us to make use of the SIMD extensions to improve throughput during the filtering step, the covariance matrix computation, the whitening, and finally the classification. Besides the throughput aspect, it helps keeping the memory footprint low. Our total memory requirement amounts to 84 kB, consisting of 2.22·875 value for the inputs and outputs of the IIR filters in 8-bit each, $18 \cdot (22+1) \cdot 22/2$ values for the W_k in 32-bit and reused for L_k , $18 \cdot (22+1) \cdot 22/2$ values for the trained feature extractor parameters $C_{\text{ref},k}^{-1/2}$ in 16-bit, and 4554.4 for the linear SVM weights in 8-bit. Compared to other embedded MI-BMI implementations, reported in Tables I and II, our model achieves the highest classification accuracy.

When using a MLP classifier, its weights need to be stored, which are larger than that of the linear SVM and dependent on its hyperparameters. This leads to a trade-off which we visualize and compare to related works in Fig. 1 by limiting the search space for the classifiers to those that let the system fit within the memory footprint indicated. This trade-off naturally crosses the implementation on Mr. Wolf [39] that was restricted to the linear SVM and expands it both to a lower as well as larger memory footprint that is only feasible on the Vega platform. It also clearly shows that our model spans the entire Pareto front within reach of typical microcontrollers as well as Mr. Wolf and Vega. The remainder of the Pareto front only includes EEG-TCNet as an additional point providing around 2% accuracy gain at the cost of a memory footprint just slightly in excess of 1.5 MB available on Vega.

B. Compute Time and Energy

Table III shows the computation time and the performance impact of the optimizations. Each output sample of the IIR filter requires 10 MACs, 3 shuffle operations, and 4 bit-shifts, resulting in a theoretical maximum of 5 MACs per cycle. Our implementation achieves 4.45 MACs per cycle with 7.24× parallel speedup compared to the single core implementation.

The covariance matrix computation reaches 10.14 MACs per cycle and a parallel speedup of 7.81× using 9 cores. The parallel speedup of the whitening is only 4.31× due to the parallelization overhead that is more visible with smaller matrix sizes (here 22×22). On the other hand, its computation accounts for 9% of the whole execution time of the final MRC, meaning that any further optimizations would only yield marginal improvements. Contrarily, the matrix logarithm is the bottleneck of the execution, requiring almost three quarters of the total number of cycles in the baseline. With the improved EVD, i.e., the optimized Householder transformation, the relative runtime is reduced to 42.58% which translates to 3.72× speedup compared to the baseline. A further improvement is given by the parallel computation which is 7.60× and 28.17× faster compared to the single core implementation with improved EVD and the baseline, respectively.

We measure the average power consumption and the execution time of all three configurations including the cluster activation and initialization, reported in Table III. The multi-core implementation consumes 2.2× more power than the single core implementations. However, the energy consumption is 6.3× and 3× lower than the baseline and the single core implementation with improved EVD, respectively, thanks to the shorter execution time.

Fig. 9 shows the energy consumption and the execution time of MRC with linear SVM measured by varying the frequency of the FC and the cluster cores of Vega. In general, increasing the clock speed yields lower energy consumption, while also lowering the runtime. With 0.6 V voltage supply, we can clock the cores up to 160 MHz. This has proven to be the most energy-efficient configuration when both FC and cluster cores run at the same frequency, consuming 210 μJ including the cluster activation, the feature extraction, and the classifier. We further fine-tune the FC frequency by clocking it at a lower speed, since all the computation is delegated to the cluster. By reducing it down to 50 MHz, the energy consumption is further reduced to 198 μJ, represented with a red cross in Fig. 9. Compared to the previous SoA low-energy embedded implementation for MI-BMI, i.e., Q-EEGNet [13], it consumes 1.7× less energy while being 3.2% more accurate. The accuracy can be further increased by 0.9% when subject-specific classifiers are trained, as discussed in Sec. VI-A. Comparing to the MRC on Mr. Wolf, it is 6.6× more energy-efficient and requires less computation time (16.90 ms vs. 33.39 ms [39]). The runtime can be further shortened by clocking the cores at higher frequency. With the supply voltage of 0.8 V, and both domains clocked at 350 MHz, we obtain an energy consumption of 338 μJ and a runtime of only 7.9 ms.

Fig. 10 depicts the measured power trace in the most energy-efficient configuration. As explained in Sec. V, the filtering, the covariance matrix, and the whitening are executed consecutively, framed with blue dashed line, while each step is computed in parallel using all cluster cores, colored with blue background. This is clearly visible from the 18 repeated patterns in the measure power trace. An additional pattern follows repeated twice. It reflects the parallel computation of the matrix logarithm, framed with red dashdotted line. The 18 matrix logarithms are computed, distributed to the 9 cores

on a first-come first-served schedule. Unlike on Mr. Wolf, where this workload is unbalanced for its 8 parallel cores and contributes negatively to the parallel speedup, the work can be evenly distributed to the 9 cores of Vega with an optimal execution. Moreover, with 4 FPU's instead of 2, our implementation reaches 2.72 floating point operations (FLOPs) per cycle compared to the 1.69 of [39]. The ideal FLOPs/cycle is not reached due to the divisions and the square root operations. Finally, the SVM accounts for a minimal part of the execution with 0.15 ms.

VII. CONCLUSION

This paper presents an improved MRC with a reduced model size while keeping comparable accuracy (75.1–76.4% vs. 75.5% [40]), allowing accurate low-power embedded BMI. We further scale down the model by quantizing and proposing a mixed-precision implementation yielding a minimal accuracy loss of 1.4%, which is still 3.7% more accurate than the SoA embedded CNN for BMI named Q-EEGNet [13]. We propose a parallel implementation on a low-power MCU called Vega, which takes only 7.9 ms and consumes 338 μ J, or alternatively takes 16.9 ms using merely 198 μ J to classify the 3.5s time window of data—57 μ W when continuously acquiring new data and processing it immediately, or 85 μ W when overlapping the 3.5s samples by 50% to avoid missing any user inputs. We provide an insight on accuracy-cost trade-off for embedded BMI models with actual implementation and measurements and claim highest accuracy implementation of EEG MI-BMI of any current work within the memory constraints of current MCUs while consuming the least energy.

REFERENCES

- [1] N. Kobayashi and M. Nakagawa, "BCI-based control of electric wheelchair using fractal characteristics of EEG," *IEEJ Tran. on Electrical and Electronic Engineering*, vol. 13, no. 12, pp. 1795–1803, 2018.
- [2] M. Vilela and L. R. Hochberg, "Chapter 8 - applications of brain-computer interfaces to the control of robotic and prosthetic arms," in *Brain-Computer Interfaces*, ser. Handbook of Clinical Neurology, N. F. Ramsey and J. del R. Millán, Eds. Elsevier, 2020, vol. 168, pp. 87–99.
- [3] K. Koizumi, K. Ueda, and M. Nakao, "Development of a cognitive brain-machine interface based on a visual imagery method," in *2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2018, pp. 1062–1065.
- [4] A. A. Frolov, O. Mokienko *et al.*, "Post-stroke rehabilitation training with a motor-imagery-based brain-computer interface (BCI)-controlled hand exoskeleton: A randomized controlled multicenter trial," *Frontiers in neuroscience*, vol. 11, p. 400, 2017.
- [5] A. Biasucci, R. Leeb *et al.*, "Brain-actuated functional electrical stimulation elicits lasting arm motor recovery after stroke," *Nature Communications*, vol. 9, no. 1, p. 2421, 12 2018.
- [6] O. Valentin, M. Ducharme *et al.*, "Validation and benchmarking of a wearable EEG acquisition platform for real-world applications," *IEEE Transactions on Biomedical Circuits and Systems*, vol. 13, no. 1, pp. 103–111, 2019.
- [7] I. Marin, M. J. H. Al-Battbooti, and N. Goga, "Drone control based on mental commands and facial expressions," in *2020 12th International Conference on Electronics, Computers and Artificial Intelligence (ECAI)*, 2020, pp. 1–4.
- [8] C. Beach, E. Balaban, and A. J. Casson, "Chapter 14 - edge algorithms for wearables: an overview of a truly multi-disciplinary problem," in *Wearable Sensors*, 2nd ed., E. Sazonov, Ed. Oxford: Academic Press, 2021, pp. 379–414.
- [9] B. Zhu, M. Farivar, and M. Shoaran, "ResOT: Resource-efficient oblique trees for neural signal classification," *IEEE Transactions on Biomedical Circuits and Systems*, vol. 14, no. 4, pp. 692–704, 2020.
- [10] D. L. T. Wong, Y. Li *et al.*, "Resource and energy efficient implementation of ECG classifier using binarized CNN for edge AI devices," in *Proc. IEEE International Symposium on Circuits and Systems (ISCAS)*, 2021.
- [11] E. Flamand, D. Rossi *et al.*, "GAP-8: A RISC-V SoC for AI at the edge of the IoT," in *2018 IEEE 29th International Conference on Application-specific Systems, Architectures and Processors (ASAP)*, 2018.
- [12] X. Wang, M. Magno *et al.*, "FANN-on-MCU: An open-source toolkit for energy-efficient neural network inference at the edge of the internet of things," *IEEE Internet of Things Journal*, 2020.
- [13] T. Schneider, X. Wang *et al.*, "Q-EEGNet: an energy-efficient 8-bit quantized parallel EEGNet implementation for edge motor-imagery brain-machine interfaces," in *2020 IEEE International Conference on Smart Computing (SMARTCOMP)*, 2020, pp. 284–289.
- [14] A. Pullini, D. Rossi *et al.*, "Mr. Wolf: An energy-precision scalable parallel ultra low power SoC for IoT edge processing," *IEEE Journal of Solid-State Circuits*, vol. 54, no. 7, pp. 1970–1981, 2019.
- [15] V. Kartsch, G. Tagliavini *et al.*, "BioWolf: A sub-10-mW 8-channel advanced brain-computer interface platform with a nine-core processor and BLE connectivity," *IEEE transactions on biomedical circuits and systems*, vol. 13, no. 5, pp. 893–906, 2019.
- [16] D. Rossi, F. Conti *et al.*, "A 1.3TOPS/W @ 32GOPS fully integrated 10-core SoC for IoT end-nodes with 1.7uW cognitive wake-up from MRAM-based state-retentive sleep mode," in *2021 IEEE International Solid-State Circuits Conference (ISSCC)*, vol. 64, 2021, pp. 60–62.
- [17] O. Russakovsky, J. Deng *et al.*, "Imagenet large scale visual recognition challenge," *International journal of computer vision*, vol. 115, no. 3, pp. 211–252, 2015.
- [18] L. Cavigelli, P. Hager, and L. Benini, "CAS-CNN: A deep convolutional neural network for image compression artifact suppression," in *IEEE International Joint Conference on Neural Networks*, 2017, pp. 752–759.
- [19] L. Deng, G. Hinton, and B. Kingsbury, "New types of deep neural network learning for speech recognition and related applications: An overview," in *2013 IEEE international conference on acoustics, speech and signal processing*, 2013, pp. 8599–8603.
- [20] T. Young, D. Hazarika *et al.*, "Recent trends in deep learning based natural language processing [review article]," *IEEE Computational Intelligence Magazine*, vol. 13, no. 3, pp. 55–75, 2018.
- [21] Q. Guo, F. Zhuang *et al.*, "A survey on knowledge graph-based recommender systems," *IEEE Transactions on Knowledge and Data Engineering*, 2020.
- [22] M. S. Seyfioglu, A. M. Özbayoğlu, and S. Z. Gürbüz, "Deep convolutional autoencoder for radar-based classification of similar aided and unaided human activities," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 54, no. 4, pp. 1709–1723, 2018.
- [23] R. T. Schirrmester, J. T. Springenberg *et al.*, "Deep learning with convolutional neural networks for EEG decoding and visualization," *Human Brain Mapping*, vol. 38, no. 11, pp. 5391–5420, 2017.
- [24] V. J. Lawhern, A. J. Solon *et al.*, "EEGNet: a compact convolutional neural network for EEG-based brain-computer interfaces," *Journal of Neural Engineering*, vol. 15, no. 5, p. 056013, 2018.
- [25] H. Wu, Y. Niu *et al.*, "A parallel multiscale filter bank convolutional neural networks for motor imagery EEG classification," *Frontiers in Neuroscience*, vol. 13, Nov. 2019.
- [26] T. M. Ingolfsson, M. Hersche *et al.*, "EEG-TCNet: An accurate temporal convolutional network for embedded motor-imagery brain-machine interfaces," in *2020 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 2020, pp. 2958–2965.
- [27] J. León, J. J. Escobar *et al.*, "Deep learning for EEG-based motor imagery classification: Accuracy-cost trade-off," *PLOS ONE*, vol. 15, no. 6, pp. 1–30, 06 2020.
- [28] F. Lotte, L. Bougrain *et al.*, "A review of classification algorithms for EEG-based brain-computer interfaces: A 10-year update," *Journal of Neural Engineering*, vol. 15, 02 2018.
- [29] K. K. Ang, Z. Y. Chin *et al.*, "Filter bank common spatial pattern (FBCSP) in brain-computer interface," in *Proc. IEEE International Joint Conference on Neural Networks*. IEEE, 2008, pp. 2390–2397.
- [30] S.-H. Park, D. Lee, and S.-G. Lee, "Filter bank regularized common spatial pattern ensemble for small sample motor imagery classification," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 26, no. 2, pp. 498–505, 2018.
- [31] C. H. Nguyen and P. Artemiadis, "EEG feature descriptors and discriminant analysis under riemannian manifold perspective," *Neurocomputing*, vol. 275, pp. 1871–1883, 2018.
- [32] F. Wu, A. Gong *et al.*, "A new subject-specific discriminative and multi-scale filter bank tangent space mapping method for recognition of

- multiclass motor imagery,” *Frontiers in Human Neuroscience*, vol. 15, p. 104, 2021.
- [33] S. Kumar, F. Yger, and F. Lotte, “Towards adaptive classification using riemannian geometry approaches in brain-computer interfaces,” in *2019 7th International Winter Conference on Brain-Computer Interface (BCI)*, 2019, pp. 1–6.
- [34] M. Congedo, A. Barachant, and R. Bhatia, “Riemannian geometry for EEG-based brain-computer interfaces: a primer and a review,” *Brain-Computer Interfaces*, vol. 4, pp. 1–20, 03 2017.
- [35] F. Yger, M. Berar, and F. Lotte, “Riemannian approaches in brain-computer interfaces: A review,” *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 25, no. 10, pp. 1753–1762, 2017.
- [36] S. Chevallier, E. Kalunga *et al.*, “Review of riemannian distances and divergences, applied to SSVEP-based BCI,” *Neuroinformatics*, 06 2020.
- [37] K. Belwafi, O. Romain *et al.*, “An embedded implementation based on adaptive filter bank for brain-computer interface systems,” *Journal of Neuroscience Methods*, 2018.
- [38] A. Malekmohammadi, H. Mohammadzade *et al.*, “An efficient hardware implementation of a motor imagery brain computer interface system,” *Scientia Iranica*, vol. 26, no. Special Issue on: Socio-Cognitive Engineering, pp. 72–94, 2019.
- [39] X. Wang, T. Schneider *et al.*, “Mixed-precision quantization and parallel implementation of multispectral riemannian classification for brain-machine interfaces,” in *2021 IEEE International Symposium on Circuits and Systems (ISCAS)*, 2021.
- [40] M. Hersche, T. Rellstab *et al.*, “Fast and accurate multiclass inference for MI-BCIs using large multiscale temporal and spectral features,” in *26th European Signal Processing Conference (EUSIPCO)*, 2018.
- [41] P. Yang, J. Wang *et al.*, “MLP with riemannian covariance for motor imagery based EEG analysis,” *IEEE Access*, vol. 8, pp. 139 974–139 982, 2020.
- [42] S. Chen, Y. Sun *et al.*, “Channel selection based similarity measurement for motor imagery classification,” in *Proc. IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, 2020, pp. 542–548.
- [43] A. Jiang, J. Shang *et al.*, “Efficient csp algorithm with spatio-temporal filtering for motor imagery classification,” *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 28, no. 4, pp. 1006–1016, 2020.
- [44] J. S. Bang, M. H. Lee *et al.*, “Spatio-spectral feature representation for motor imagery classification using convolutional neural networks,” *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–12, 2021.
- [45] Y. Li, L. Guo *et al.*, “A temporal-spectral-based squeeze-and- excitation feature fusion network for motor imagery eeg decoding,” *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 29, pp. 1534–1545, 2021.
- [46] C. Brunner, R. Leeb *et al.*, “BCI competition 2008 - Graz data set A,” <http://bnci-horizon-2020.eu/database/data-sets>.
- [47] X. Wang, M. Hersche *et al.*, “An accurate EEGNet-based motor-imagery brain-computer interface for low-power edge computing,” in *IEEE International Symposium on Medical Measurements and Applications (MeMeA)*, 2020.
- [48] A. Barachant, S. Bonnet *et al.*, “Classification of covariance matrices using a riemannian-based kernel for bci applications,” *Neurocomputing*, vol. 112, pp. 172–178, 2013, advances in artificial neural networks, machine learning, and computational intelligence.
- [49] B. Jacob, S. Kligys *et al.*, “Quantization and training of neural networks for efficient integer-arithmetic-only inference,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 2704–2713.
- [50] L. Cavigelli and L. Benini, “RPR: random partition relaxation for training; binary and ternary weight neural networks,” *CoRR*, vol. abs/2001.01091, 2020.
- [51] M. Spallanzani, L. Cavigelli *et al.*, “Additive noise annealing and approximation properties of quantized neural networks,” *CoRR*, vol. abs/1905.10452, 2019.
- [52] J. O. Smith, *Introduction to Digital Filters with Audio Applications*. W3K Publishing, 2007, accessed on: Nov. 8, 2020. [Online] Available: <http://ccrma.stanford.edu/jos/filters/>.
- [53] J. H. Wilkinson, “Error analysis of eigenvalue techniques based on orthogonal transformations,” *Journal of the Society for Industrial and Applied Mathematics*, vol. 10, no. 1, pp. 162–195, 1962.
- [54] X. Wang, “DSP library for PULP,” <https://github.com/pulp-platform/pulp-dsp>, 2019.
- [55] R. Burden and J. Faires, *Numerical analysis*. Cengage Learning, 2004.
- [56] J. H. Wilkinson, *The algebraic eigenvalue problem*. Oxford Clarendon, 1965, vol. 662.
- [57] W. Givens, “Numerical computation of the characteristic values of a real symmetric matrix,” Oak Ridge National Lab., Tech. Rep., 1954.
- [58] D. Bindel, J. Demmel *et al.*, “On computing givens rotations reliably and efficiently,” *ACM Transactions on Mathematical Software (TOMS)*, vol. 28, no. 2, pp. 206–238, 2002.
- [59] A. Garofalo, M. Rusci *et al.*, “PULP-NN: accelerating quantized neural networks on parallel ultra-low-power RISC-V processors,” *Philosophical Transactions of the Royal Society A*, vol. 378, no. 2164, 2020.




Xiaying Wang received her B.Sc. and M.Sc. degrees in biomedical engineering from Politecnico di Milano, Italy and ETH Zürich, Switzerland in 2016 and 2018, respectively. She is currently pursuing a Ph.D. degree at the Integrated Systems Laboratory at ETH Zürich. Her research interests include biosignal processing, brain-machine interface, low power embedded systems, energy-efficient smart sensors, and machine learning on microcontrollers. She received the excellent paper award at the IEEE Healthcom conference in 2018 and she won the Ph.D. Fellowship funded by the Swiss Data Science Center in 2019.



Lukas Cavigelli received the B.Sc., M.Sc., and Ph.D. degree in electrical engineering and information technology from ETH Zürich, Zürich, Switzerland in 2012, 2014 and 2019, respectively. After spending a year as a Postdoc at ETH Zürich, he has joined Huawei’s Zurich Research Center in Spring 2020. His research interests include deep learning, computer vision, embedded systems, and low-power integrated circuit design. He has received the best paper award at the VLSI-SoC and the ICDESC conferences in 2013 and 2017, the best student paper award at the Security+Defense conference in 2016, and the Donald O. Pederson best paper award (IEEE TCAD) in 2019.



Tibor Schneider received his B.Sc and M.Sc in electrical engineering and Information Technology at Hochschule für Technik Rapperswil and ETH Zurich, Switzerland in 2017 and 2021, respectively. Since 2021, he is a Ph.D student at the Networked Systems Group at ETH Zürich, advised by Laurent Vanbever. His current research is aimed at improving verification and synthesis methods for configurations of internet networks. In 2021, he won the ABB Research Prize at ETH Zürich.



Luca Benini is the Chair of Digital Circuits and Systems at ETH Zürich and a Full Professor at the University of Bologna. He has served as Chief Architect for the Platform2012 in STMicroelectronics, Grenoble. Dr. Benini’s research interests are in energy-efficient system and multi-core SoC design. He is also active in the area of energy-efficient smart sensors and sensor networks. He has published more than 1’000 papers in peer-reviewed international journals and conferences, four books and several book chapters. He is a Fellow of the ACM and of the IEEE and a member of the Academia Europaea.