

Embedded Low-Power Processor for Personalized Stress Detection

Nasrin Attaran¹, Abhilash Puranik¹, Justin Brooks², and Tinoosh Mohsenin¹

¹Department of Computer Science & Electrical Engineering, University of Maryland, Baltimore County

²Human Research and Engineering Directorate, US Army Research Lab

Abstract—Personal monitoring systems require sampling and processing on multiple streams of physiological signals to extract meaningful information. These systems require a large number of digital signal processing and machine learning kernels which typically require significant amounts of power. However, to be used in a wearable environment, the processing system needs to be low-power, real-time, and light-weight. In this paper, we present a personalized stress monitoring processor that can meet these requirements. First, various physiological features are explored to maximize stress detection accuracy using two machine learning classifiers including Support Vector Machine (SVM) and K-Nearest Neighbors (KNN). Among different extracted features from four physiological sensors, heart rate and accelerometer features have 96.7% (SVM) and 95.8% (KNN) detection accuracy. In the second part, two fully flexible and multi-modal processing hardware designs are presented that consist of feature extraction and classification algorithms. We first demonstrate the ASIC post-layout implementation of both designs in 65 nm CMOS technology as well as the implementation on Artix-7 FPGA. The proposed SVM and KNN processors on the ASIC platform occupy an area of 0.17 mm² and 0.3 mm² and dissipate 39.4 mW and 76.69 mW power, respectively. The ASIC implementation improves the energy efficiency by 42x (SVM) and 12x (KNN) over FPGA implementations. The entire stress monitoring system is further evaluated against a number of other platforms including Raspberry Pi 3B, NVIDIA TX1 GPU and NVIDIA TX2 GPU. The experimental results indicate that ASIC and FPGA platforms have the highest throughput (decision/sec) as well as lowest power consumption over all other platforms. The ASIC/FPGA implementations improve the energy efficiency (throughput/power) by 6/5 and 5/4 order of magnitude compared to TX1 GPU and Raspberry pie ARM platforms, respectively.

Index Terms—ASIC, FPGA, K-nearest neighbor (KNN), personalized stress detection, support vector machine (SVM)

I. INTRODUCTION

Personalized wearable biomedical systems enable the acquisition of various physiological and behavioral data that can be used to make general inferences about the state of human [1]. These systems need to process parallel streams of multi-physiological data in real-time and within a limited power budget. Furthermore, they often utilize applications which require a large number of digital signal processing (DSP) and machine learning (ML) techniques to extract meaningful knowledge. Stress detection is one such health monitoring application. In fact, there have been several recent studies on stress detection that have used multiple physiological signals, such as electrocardiogram (ECG), heart rate (HR), electrodermal activity (EDA), electromyography (EMG), respiration, blood pressure and oxygen saturation (SpO₂) [2], [3], [4], [5].

The relatively high amount of power consumption and delays required to transmit raw or even compressed data, make it essential to process sensor data locally on-board [6], [7], [8], [9], [10], [11], [12], [13]. For this class of personalized biomedical applications, the sampling frequency is relatively low in the range of 50 Hz to 2 KHz. Thus most of the current processing platforms can meet the sampling deadlines when running at high clock frequency. However, they cannot fit within the stringent power budget of wearable devices.

The main focus of this work is designing a low-power, lightweight and real-time local processor for a personalized stress detection system. All design steps can be used and modified for other biomedical applications based on their specific characteristics including biomedical signals, computational latency and energy consumption. All sensor data is processed locally on-board, and only a final decision is sent out, rather than sending preprocessed raw sensor data. For this purpose, we examined the implementation of stress detection on ASIC and FPGA platforms as well as embedded commercial-off-the-shelf platforms including Raspberry Pi 3B, NVIDIA Jetson TX1 GPU, and NVIDIA Jetson TX2 GPU. The main contributions of the paper include:

- Proposing an efficient local processor for stress detection using multi-modal physiological signals, feature extraction techniques and machine learning classifiers.
- Thorough analysis of features extracted from multi-sensor data with SVM and KNN ML classifiers to find the best feature sets and their combination.
- Implementing the flexible on-board processor on ASIC and FPGA platforms for personalized stress detection.
- Comparing the stress detection processor's performance metrics with existing embedded commercial-off-the-shelf platforms including ARM A54 CPU on Raspberry Pi and NVIDIA GPUs on Jetson TX1 and TX2.

The remainder of this paper describes the different sections of multi-modal stress monitoring system, two different reconfigurable hardware designs, hardware experimental results and comparison with embedded off-the-shelf platforms.

II. CLASSIFIER AND FEATURE EXTRACTION ANALYSIS

A. Description of Dataset

In this paper, we used the data from a naturalistic shooting task which consists of multi-physiological and behavioral recordings from 15 subjects [14].




Figure 1. Block diagram of a multi-physiological stress detection system containing data acquisition by sensors, feature extraction, and machine learning classifier to generate the result.

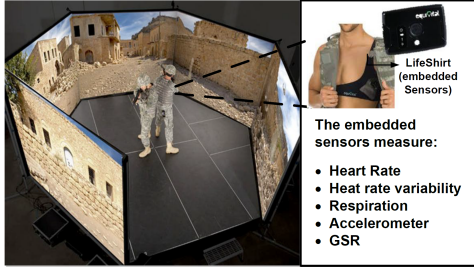


Figure 2. 300-degree simulator to collect the multi-physiological data during different levels of stress using the embedded sensors in wearable LifeShirt [14].

The participants performed a shooting task in a simulator in which they had to discriminate enemy versus friendly targets and decide to shoot or refrain respectively. Three levels of stress were induced by manipulating performance feedback on incorrect trials in a blocked design: low (No feedback), medium (visually displayed or Lifebar), and high (Shock). This dataset was collected for one experiment to evaluate the effect of different levels of stress on participants' behavior. It has various physiological recordings and can be used for stress detection evaluation. For this paper, only the low-stress and high-stress conditions were studied.

B. Feature Extraction

Using the raw physiological data in health monitoring systems, which need to process multiple signals, requires a significant amount of memory and power. Feature extraction reduces the number of resources necessary to describe a large set of data. Finally, using the extracted features rather than raw data improves the classification and clustering accuracy.

Figure 3 shows raw data representation of HR and accelerometer signals for individual 13. The levels of accelerometer and HR signals are entirely distinctive in shock and nofeedback condition. A total of 16 features used in the experiment were derived in 35-second windows and are elaborated in Table I. Figure 4 represents mean HR and mean ACC.X features from individual 13.

C. Classifier Selection and Feature Selection

After feature extraction, we need to examine which features contain the most useful information and remove those features that do not improve the model. In this work, we used classification accuracy as an automated method to choose the most appropriate features. To find the best combination of the features, we examined the classification accuracy for

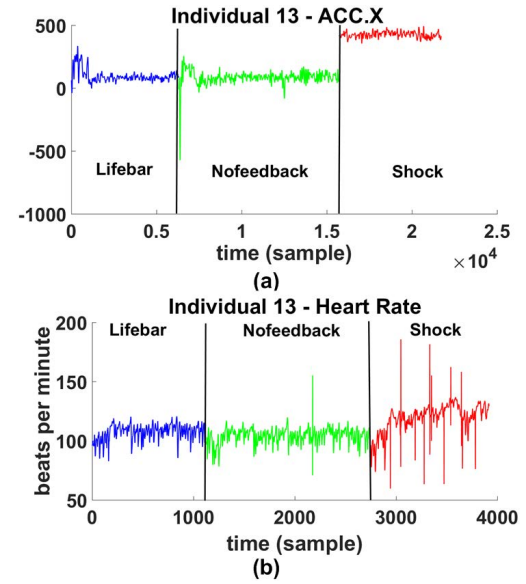


Figure 3. Raw signals' representation for participant 13. (a) Accelerometer (X axis) and (b) HR.

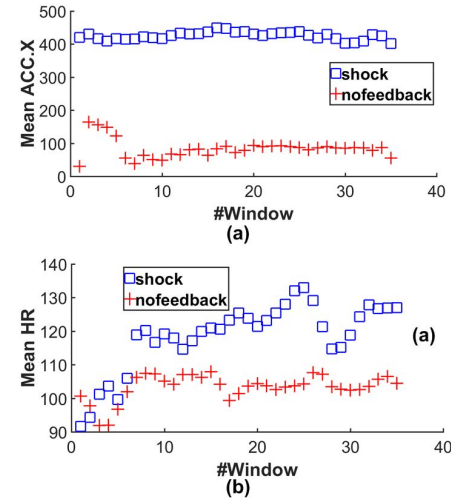


Figure 4. Extracted features representation for participant 13. (a) mean Axis-x and (b) mean HR.

each feature for all individuals independently. We utilized two popular supervised binary classifiers, SVM and KNN. Figure 5 shows the average accuracy across all 15 individuals for each feature for both SVM and KNN classifiers. The seven highlighted features achieve the highest accuracy in stress detection across all individuals. Since our final goal is to implement one low-power and lightweight processor for stress monitoring, we are interested in finding the best reduced feature set to detect stress accurately. We utilized attribute selection function in WEKA tool to find the best feature set among these seven features. The final feature set has four features including mean HR, mean Acc.X, mean Acc.Y and mean Acc.Z. Table II shows the average accuracy of each feature and the concatenated feature set. For SVM and KNN classifiers, the detection accuracy of feature set is on average 17.94% and 10.95%, higher than that of individual features, respectively. Thus, we designed and proposed the following hardware processors for stress detection based on the concatenated feature set.

Feature No.	Sensors	Features
1 to 7	ECG	Mean HR, Std HR, Mean RR, Std RR LF-HRV, HF-HRV, LF/HF ratio
8 to 14	Accelerometer	Mean of X,Y and Z axis Standard deviation of X,Y and Z axis Magnitude of three axes
15	Resp. Rate	Mean RR
16	SpO ₂	Mean SpO ₂

Table I

16 EXTRACTED FEATURES FROM FOUR PHYSIOLOGICAL SENSORS PER EACH 35-SECOND WINDOW.




Figure 5. Average of two-level stress detection accuracy using multi-physiological sensors and corresponding features for 15 individuals. The features with the highest accuracy (more than 65% for both SVM and KNN) are highlighted.

III. FEATURE EXTRACTION AND RECONFIGURABLE CLASSIFIER HARDWARE

A. Proposed Scalable and Pipeline SVM Processor

Figure 6.a shows a detailed architecture of the linear SVM processor used for a stress monitoring system based on four extracted features from the heart rate and accelerometer sensors. The proposed parallel pipelined architecture can be easily reconfigured to process any number of features and support vectors for a variety of applications. The support vectors (SV), bias (b) and other required coefficients were calculated offline using the SVMtrain MATLAB function. Each memory block is loaded with pre-computed weighted support vectors from a trained model for each feature. There are sufficient registers in this design to store the intermediate results in the pipeline scheme. The classifier receives the features derived in 35-second windows as testing input. The dot product operation runs between the testing data and all supporting vectors available in RAM blocks. This is followed by the parallel dot product operation, which is added with bias parameter at the final stage to find the prediction result.

B. Proposed Scalable Semi-Parallel KNN Processor

Figure 6.b shows a high-level block diagram of the KNN processor used for the stress monitoring system based on four extracted features from the heart rate and accelerometer sensors and can be reconfigured to process any number of feature vectors for different applications. The training samples (256 training samples for personalized stress monitoring) and their corresponding labels are stored in the ROM block. The extracted feature set (4 features) from testing samples from the given 35-second window is stored in the buffer component. The four subtractors, four multipliers, and one adder modules are used to find the Euclidean distance between the given test sample and training data in parallel. However, the training data was read from the ROM block serially. The sorting block is responsible for finding the K smallest distance between

Classifier	SVM	KNN
Mean HR	78.99%	83.06%
Mean Acc.X	88.79%	92.13%
Mean Acc.Y	72.83%	76.66%
Mean Acc.Z	74.38%	87.56%
Combined feature set	96.7%	95.81%

Table II

THE AVERAGE ACCURACY OF THE FOUR MOST IMPORTANT FEATURES AS WELL AS CONCATENATED FEATURE SET ACROSS ALL 15 PARTICIPANTS. FOR SVM AND KNN CLASSIFIERS THE DETECTION ACCURACY OF CONCATENATED FEATURE SET ON AVERAGE 17.94% AND 10.95% HIGHER THAN THAT OF INDIVIDUAL FEATURES, RESPECTIVELY.

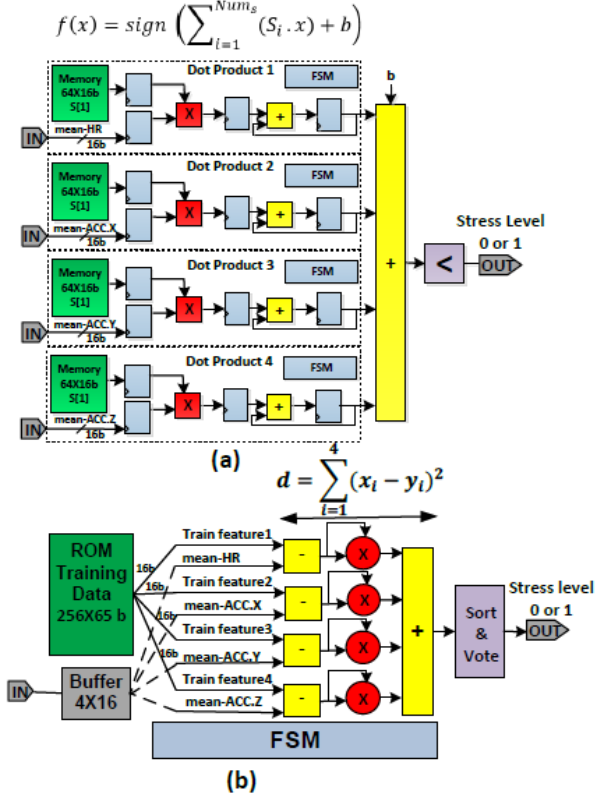


Figure 6. Block diagram of the proposed reconfigurable processors using (a) SVM and (b) KNN for personalized stress detection system. The concatenated feature set mean-HR, mean ACCX, mean-ACCY and mean ACCZ are stored in the memory blocks.

the testing sample with all training data. The voting module generates the label of the testing sample based on the majority voting at the final step. The Finite state machine (FSM) module is responsible for syncing and controlling all components in the design. This design is fully configurable for a variable number of features and different size of training data.

IV. ASIC AND FPGA IMPLEMENTATION AND RESULTS

A. ASIC Results

The stress detection processors for both SVM and KNN classifier configurations are synthesized and placed and routed in the 65 nm TSMC CMOS technology. Figures 7 and 8 shows the layout of the proposed stress detection processors (features extraction + classifier) and post layout results.

The SVM processor occupies 0.17 mm² and dissipates approximately 39.4 mW when running at its maximum frequency of 250 MHz. When the chip operates at the nominal frequency of 5 Hz required to meet the 17.5-second deadline, it dissipates 0.76 nW (linearly scaled with frequency), which results in 13.4 nJ at 1 V to classify one 35-second window




Figure 7. Layout view and post-layout implementation results of the proposed multi-modal SVM processor (64 support vectors) + feature extraction. The highlighted regions indicate the location of four dot product components and feature extraction on the chip. *The power and energy are reported for the nominal frequency where the computation is done in 17.5-second interval window.




Figure 8. Layout view and post-layout implementation results of the proposed multi-modal KNN processor (256 training data) + feature extraction. The highlighted regions indicate the location of training memory, sorting, distance calculation and feature extraction on the chip. *The power and energy are reported for the nominal frequency where the computation is done in 17.5-second interval window.

of input. The KNN processor runs at the nominal frequency of 59 Hz and dissipates the power of 17.96 nW. It consumes 0.31 μJ for stress detection per window.

B. FPGA Results

As a second hardware-based platform for stress detection implementation, we utilized Xilinx Artix-7 FPGA. FPGAs are highly flexible allowing on-the-fly configuration to optimize bit resolution, clock frequency, parallelization, and pipelining for a given application. The main disadvantages of FPGAs, however, are that they have substantially higher leakage power and require writing low-level logic blocks in HDL [6]. For the stress detection case study, complete FPGA hardware for SVM and KNN machine learning kernels in addition to feature extraction were developed in Verilog that utilized highly parallel, highly pipelined DSP and ML kernels. Both real-time and simulated projections using commercial tools were used to perform timing and power analysis when running test stimulus. For stress detection application, the Artix-100T FPGA is targeted on the Nexys platform. Table III summarizes the results of implementing the stress detection case study using SVM and KNN on Artix-7 FPGA.

V. SOFTWARE-BASED PLATFORMS

A. NVIDIA Jetson TK1

B. Raspberry Pi

VI. COMPARISON WITH EMBEDDED OFF-THE-SHELF PLATFORMS

We examined the implementation of stress detection on several embedded commercial-off-the-shelf processors including Raspberry Pi 3B, Jetson TX1 GPU, and Jetson TX2 GPU as well as hardware platforms.

Design	SVM	KNN	Improvement
Registers (#)	278	440	1.6x
LUTs (#)	197	692	3.5x
Memory (Kb)	4	16	4x
Max Freq. (MHz)	200	200	-
Latency (cycles)	170	1025	6x
Latency (us)	0.85	5.125	6x
Nominal Freq. (Hz)	58	9	6.5
Dynamic Power (nW) ¹	4.95	37.48	7.6x
Leakage Power (mW)	82	82	-
Energy (μJ) ²	0.08	0.65	8x

Table III

THE COMPARISON OF STRESS DETECTION HARDWARE IMPLEMENTATION (CLASSIFIER + FEATURE EXTRACTION) FOR SVM AND KNN CLASSIFIER ON ARTIX-7 FPGA. 1. THE DYNAMIC POWER RESULTS ARE FOR THE NOMINAL FREQUENCY TO MEET THE 17.5-SECOND WINDOW INTERVAL. 2. SINCE FPGA HAS SIGNIFICANT LEAKAGE POWER (DOMINATE COMPARED TO DYNAMIC POWER), THE ENERGY RESULTS ARE BASED ON DYNAMIC POWER.




Figure 9. Comparison of energy-delay-product (EDP) for the stress detection case study with KNN and SVM classifiers when implemented on several processor combinations including Raspberry Pi, Jetson TX1 GPU, Jetson TX2 GPU, Artix-7 FPGA and ASIC. The Raspberry Pi was considered as a baseline. The ASIC implementations for both designs have the lowest EDP.

Raspberry Pi 3B is packed with 1.2 GHz Quad-core ARMv8 CPU and 1GB LPDDR2 RAM. It features Bluetooth and Wireless connectivity and powered by the Broadcom BCM2837 SoC.

Jetson TX1 and Jetson TX2 development boards are packed with a 256-core NVIDIA Maxwell-based GPU and NVIDIA Pascal-based GPU, respectively. We used the serial C code to perform stress detection on Raspberry Pi single CPU core. For data-level parallelization execution on the GPU, we used PyCuda. We utilized only one block on GPU to parallelize both KNN (with 256 threads for 256 training data) and SVM (64 threads for 64 support vectors). For a fair comparison among different platforms, the power of Nexys board is added to FPGA results. Tables IV and V show the comparison results for all platforms for both classifiers.

The ASIC is the best among all other platforms with respect to throughput, energy consumption, and energy efficiency. To better understand the benefit of ASIC and FPGA implementation of stress detection, Figure 9 provides energy-delay-product (EDP) comparison among all platforms. The ASIC implementation has significantly lower EDP than all other platforms. Minimizing the EDP is essential for biomedical applications, as it is critical to both promptly making decisions and consuming minimal energy. The ASIC has 16x and 100x lower EDP compared to FPGA for KNN and SVM implementation, respectively. Furthermore, ASIC's energy efficiency is 11x and 42x larger than FPGA for KNN and SVM, respectively. Although ASIC hardware implementation significantly improves energy efficiency, but it may not be practical due to the cost and time to market constraints. The FPGA

Processor	Clock (MHz)	Power (mW)	Throughput (dec/sec)	Energy (mJ)	Energy Efficiency (dec/sec/watt)	Energy Efficiency Improvement (over baseline)
ARM A53 (baseline)	900	1,480	2	746.36	1.33	1x
TX2 GPU	854	2,120	130.54	16.23	61.58	46x
TX1 GPU	998	2,430	225	10.76	92.89	69x
Artix-7 100T FPGA	200	728	195,121	0.0035	268,024	200,044x
ASIC	250	76.69	243,902	0.0003	3,180,368	2,373,712x

Table IV

BREAKDOWN OF HARDWARE RESULTS FROM RUNNING STRESS DETECTION APPLICATIONS ON A VARIETY OF PROCESSING PLATFORMS (FEATURE EXTRACTION + KNN CLASSIFIER WITH 256 TRAINING SAMPLES). RESULTS INCLUDE THROUGHPUT, ENERGY, AND ENERGY EFFICIENCY. IMPLEMENTATION ON ARM A53 CPU ON RASPBERRY PI IS FULLY SERIAL ON A SINGLE CPU AND IS USED AS BASELINE FOR COMPARISON. FOR FPGA THE POWER OF NEXYS BOARD IS ADDED.

Processor	Clock (MHz)	Power (mW)	Throughput (dec/sec)	Energy (mJ)	Energy Efficiency (dec/sec/watt)	Energy Efficiency Improvement (over baseline)
ARM A53 (baseline)	900	1530	5.29	289.17	3.45	1x
TX2 GPU	854	2090	212.76	9.82	101.8	29x
TX1 GPU	998	2610	357.14	7.308	136.83	39x
Artix-7 100T FPGA	200	702	1,250,000	0.00056	1,780,626	514,903x
ASIC	250	39.4	2,941,176	0.000013	74,649,149	21,586,294x

Table V

BREAKDOWN OF HARDWARE RESULTS FROM RUNNING STRESS DETECTION APPLICATIONS ON A VARIETY OF PROCESSING PLATFORMS (FEATURE EXTRACTION + SVM CLASSIFIER WITH 64 SUPPORT VECTORS). RESULTS INCLUDE THROUGHPUT, ENERGY, AND ENERGY EFFICIENCY. IMPLEMENTATION ON ARM A53 CPU ON RASPBERRY PI IS FULLY SERIAL ON A SINGLE CPU AND IS USED AS BASELINE FOR COMPARISON. FOR FPGA THE POWER OF NEXYS BOARD IS ADDED.

solution achieves the second EDP for stress detection and offers reprogrammability and low development cost compared to the ASIC implementation.

VII. CONCLUSION

Health monitoring applications share strong commonalities, including requiring sampling from several physiological signals at various rates, preprocessing, feature extraction and machine learning kernels. In this paper, we demonstrated an accurate stress monitoring system by utilizing multiple physiological signals. Our analysis indicated that using heart rate and accelerometer signals for determining the level of stress, generated the most accurate classification with both KNN and SVM classifiers. The average accuracy of the personalized stress monitoring system with KNN and SVM classifiers are 95.8% and 96.7%, respectively. This research also examined the choice of various processors including ASIC, FPGA, Raspberry Pi, NVIDIA TX1 and TX2 for energy-efficient processing of physiological signals for the multi-modal stress detection application. The experimental results showed that the post-layout (ASIC) implementation of the SVM and KNN processors minimizes power consumption and latency as well as maintaining a low-area footprint for personalized stress monitoring. The ASIC implementation improves the energy efficiency by 42x and 12x over FPGA platform for SVM and KNN implementations, respectively.

VIII. ACKNOWLEDGEMENT

The authors would like to thank Debbie Patton and Justin Brooks at Army research laboratory for providing the dataset in this work. This research is based upon work partially supported by the National Science Foundation under Grant No. 00010145 to develop processors for biomedical applications.

REFERENCES

- [1] N. Attaran, J. Brooks, and T. Mohsenin, "A low-power multi-physiological monitoring processor for stress detection," in *2016 IEEE SENSORS*, Oct 2016, pp. 1–3.
- [2] J. A. Healey *et al.*, "Detecting stress during real-world driving tasks using physiological sensors," *IEEE Transactions on Intelligent Transportation Systems*, vol. 6, no. 2, pp. 156–166, June 2005.
- [3] Y. Shi, M. H. Nguyen, P. Blitz, B. French, S. Fisk, F. De la Torre, A. Smailagic, D. P. Siewiorek, M. al'Absi, E. Ertin *et al.*, "Personalized stress detection from physiological measurements," in *International symposium on quality of life technology*, 2010, pp. 28–29.
- [4] J. Zhai, A. B. Barreto, C. Chin, and C. Li, "Realization of stress detection using psychophysiological signals for improvement of human-computer interactions," in *SoutheastCon, 2005. Proceedings. IEEE*. IEEE, 2005, pp. 415–420.
- [5] F.-T. Sun *et al.*, "Activity-aware mental stress detection using physiological sensors," in *International Conference on Mobile Computing, Applications, and Services*. Springer, 2010, pp. 211–230.
- [6] A. Page *et al.*, "Low-power manycore accelerator for personalized biomedical applications," in *Proceedings of the 26th Edition on Great Lakes Symposium on VLSI*, ser. GLSVLSI '16. New York, NY, USA: ACM, 2016, pp. 63–68.
- [7] S. Viseh, M. Ghovanloo, and T. Mohsenin, "Towards an ultra low power on-board processor for tongue drive system," *Circuits and Systems II: IEEE Transactions on*, accepted, vol. 62, no. 2, pp. 174–178, Feb 2015.
- [8] A. Jafari, N. Buswell, M. Ghovanloo, and T. Mohsenin, "A low-power wearable stand-alone tongue drive system for people with severe disabilities," *IEEE Transactions on Biomedical Circuits and Systems*, vol. 12, no. 1, pp. 58–67, Feb 2018.
- [9] A. Page *et al.*, "A flexible multichannel eeg feature extractor and classifier for seizure detection," *Circuits and Systems II: Express Briefs, IEEE Transactions on*, vol. 62, no. 2, pp. 109–113, 2015.
- [10] J. Han, Y. Zhang, S. Huang, M. Chen, and X. Zeng, "An area-efficient error-resilient ultralow-power subthreshold ecg processor," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 63, no. 10, pp. 984–988, Oct 2016.
- [11] T. Tekeste, H. Saleh, B. Mohammad, A. Khandoker, and M. Elnaggar, "A nano-watt ecg feature extraction engine in 65nm technology," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. PP, no. 99, pp. 1–1, 2017.
- [12] J. Lee, S. Park, I. Hong, and H. J. Yoo, "An energy-efficient speech-extraction processor for robust user speech recognition in mobile head-mounted display systems," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 64, no. 4, pp. 457–461, April 2017.
- [13] A. Kulkarni, A. Page, N. Attaran, A. Jafari, M. Malik, H. Homayoun, and T. Mohsenin, "An energy-efficient programmable manycore accelerator for personalized biomedical applications," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. PP, no. 99, pp. 1–14, 2017.
- [14] D. Patton, "How good is real enough? 300 degree of virtual immersion," *Masters Thesis, Towson University Department of Psychology*, 2013.