



An Investigation of Single-Core and Multi-Core Computing Methods for Biosignal Processing

¹Dr BALASUBRAMANIAN B, ²R SANTHIYA, ³K. ROKINI

*Department of Biomedical Engineering, Excel Engineering College (AUTONOMOUS)
Komarapalayam, Namakkal, Tamil Nadu.*

⁴S.MAHESWARI, *Department of Biomedical Engineering, Velalar College of Engineering
&Technology, Thindal, Erode -12.*

Article History	Abstract
Received: 26 June 2023 Revised: 15 August 2023 Accepted: 21 August 2023	<i>This paper provides a single-core and multi-core processor design for applications involving highly parallel processing and sluggish biosignal events in health surveillance systems. An instruction memory (IM), a data memory (DM), and a processor core (PC) make up the single-core design. In contrast, the multi-core architecture is made up of PCs, separate IMs for each core, a shared DM, and an interconnection cross-bar connecting the cores and the DM. The power vs. performance compromises for a multi-lead ECG signal conditioning application that takes advantage of near threshold computing are evaluated between both designs. According to the findings, the multi-core system uses 10.4% more power for low processing demands (681 kOps/s) and 66% less power for high processing needs (50.1 MOps/s).</i>
CC License	Keywords: <i>ECG, Computing, WBSN, Near Threshold, Parallel Processing.</i>

Introduction

In order to identify and assess any health issues, personal health systems track metabolic processes such as carbon dioxide levels and blood oxygen, cardiac and lung rates, and blood pressure. The technology that makes these personal health systems possible is called wireless body sensor networks, or WBSNs [1]. A wireless body-network (WBSN) for health surveillance is made up of several small sensor nodes that are affixed to the body and are each in charge of processing a distinct low-rate physiological signal. The electrocardiogram (ECG), for example, is one of the most significant signals generated by the body. It is usually recorded at sample rates ranging from 125 Hz to 1 kHz in order to obtain the frequently significant waveform details.

In order to identify and assess any health issues, personal health systems track metabolic processes such as carbon dioxide levels and blood oxygen, cardiac and lung rates, and blood pressure. The technology that makes these personal health systems possible is called wireless body sensor networks, or WBSNs [1]. A wireless body-network (WBSN) for health surveillance is made up of

several small sensor nodes that are affixed to the body and are each in charge of processing a distinct low-rate physiological signal. The electrocardiogram (ECG), for example, is one of the most significant signals generated by the body. It is usually recorded at sample rates ranging from 125 Hz to 1 kHz in order to obtain the frequently significant waveform details.

Unfortunately, since devices are expected to run on a single battery for extended periods of time, achieving low-power consumption remains a significant challenge, even with the reasonable required processing effort. Supply voltage scaling, possibly all the way to sub-threshold functioning, is an efficient strategy to cut power consumption. They have been thoroughly examined, and a number of low-power architectures have been offered. For instance, it suggested the use of solar cell harvesting to create a sensor platform that could run almost continuously. An ARM Cortex M3 core with both non-retentive and retentive SRAM, as well as a power control unit that regulates the active and ultra low power sleep methods, are features of the suggested single processor structure.

[1] In order to support wireless monitoring systems, this introduced a new ultra low energy CPU with low voltage operations. With the use of power gating, memory capacity and operating set modifications, and a new low permeability memory, they were able to optimise the processor's standby power usage. [2] The primary problem with low-voltage operation, however, is loss of performance, which can restrict the amount of voltage-scaling that can be done for a given processing need. If the algorithms to be conducted can be parallelized, [3] this problem can be mitigated by using many cores for parallel computation.

Suggested a cluster-based, near threshold computing (NTC) multi-processor architecture that can service several cores simultaneously. [44] It has a shared cache and runs at a higher supply voltage. In a different work, Yu et al. [5] used architectural level parallelism to compensate for the performance loss and proposed a sub/near threshold co-processor for energy-efficient mobile processing of images. [6] Lastly, a massively parallel stream processor working in NTC that can accomplish one giga-operation per second while consuming one milliwatt of power overall was proposed.

Application of ECG Signal Conditioning

[12] ECG involves the examination of electrical changes that occur with each heartbeat as the heart muscle depolarizes, as detected by electrodes linked to the body. [7] With a single-lead electrocardiogram (ECG), the voltage difference between two electrodes positioned on either side of the heart reveals heart rate and makes it possible to spot deficiencies in various heart muscle sections. [8] Up to 12 leads can be used to get a clearer and more comprehensive image of the heart muscle activity. [9] Every lead presents the heart's activity from a unique angle. Unfortunately, baseline drifts and other forms of noise are present in raw ECG data even when they are captured in a controlled environment. [10] Therefore, one of the primary uses of a sensor network in WBSNs for robotic ECG analysis or signal reduction for recording is ECG signal conditioning [11]. Thus, the ECG signal conditioning technique based on filtering morphology provided in [13] serves as our benchmark solution. This method works on many leads in parallel and independently, correcting baselines and suppressing noise in ECG readings. Unfortunately, baseline drifts and other forms of noise are present in raw ECG data even when they are captured in a controlled environment.

Therefore, one of the primary uses of a sensor network in WBSNs for robotic ECG analysis or signal reduction for recording is ECG signal conditioning [14]. Thus, the ECG signal conditioning technique based on filtering morphology provided in [15] serves as our benchmark solution. This method works on many leads in parallel and independently, correcting baselines and suppressing noise in ECG readings.

Architecture of Processing Platforms

In order to concentrate on comparing the single-core and multi-core configuration, we use the same processing unit (PU) and data memory (DM) to construct both reference designs. The ideas are put into practise using a 90 nm low leakage process technology that trades peak performance for a notable reduction in leakage power, particularly in the memories.

Processing Unit:

A processor unit (PU) consists of a processing core (PC) and a 24-bit wide instruction memory (IM) with a capacity of 4k instruction words (12 kBytes). This is adequate for numerous common biomedical applications on wireless broadband network systems, including data compression and partitioning [16], [17]. The PC has a Harvard memory model and a 16-bit architecture similar to a Reduced Instruction Set Computer (RISC), with sixteen functional registers. The application's low to moderate efficiency needs are met by the straightforward two-stage pipeline, which also lowers the number of registers that require clocking. Instructions are decoded and read addresses are generated in the first pipeline stage. The operations are carried out in the second pipeline stage, and the outcomes are saved in a working register or a data memory location. To support the energy-efficient execution of signal processing algorithms, the instruction set includes, among other things, [18] single-cycle multiplications, multi-bit shifts, and logic and arithmetic operations. The majority of instructions are performed in a single clock cycle with a two-cycle latency, occupying only one (24-bit) instruction-word.

Data Memory:

The PC has simultaneous clock cycle access to the DM for both writing and reading. As a result, the DM needs two different access ports: one for writing and one for reading. In order to process an 8-lead ECG in real time, 64 kBytes of DM must be divided into M memory banks (MBs), or 16 memory banks, with 2k words in each bank. This arrangement allows for partial shutdown for power leakage minimization for applications that have lower memory demands, and it matches the maximum output of our 2-port memory generator.

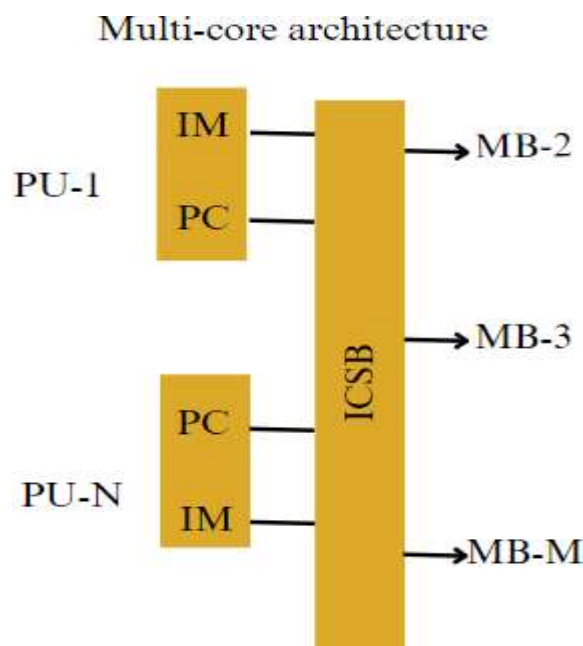


Fig. 1. Platforms for processing

Architecture of Single-Core Processors

Fig. 1(a) displays the reference architecture for a single-core WBSN sensor node. The data is multiplexed and the single PU is connected to each unique MB using a straightforward selection logic (SL). The system uses 5580 cycles per sample to process the 8-lead ECG signals consecutively. Our single-core design, when speed-optimized, could run at nominal 1.2 V supply voltage up to 147MHz, which is far more than needed for most biomedical signal processing uses, even at very high sampling rates. We optimise the design used as a reference for lowest area instead of maximum speed because the cost of this highspeed is lower energy efficiency, which lowers the active and leakage power consumption. In order to do this, we ask the EDA design tools to select logic gates that have low driving power.

The distribution of the area-optimized design's power consumption at 1.2V and 16 MHz clock frequency is displayed in Table 1's second column. It's interesting to observe that roughly 15% (0.55 mW) of the overall power consumption is used by the SL and the interconnect network (routing method and buffering) between the PU and MBs. A more thorough investigation reveals that issues with the address and data buses are mostly to blame for this power. We install 48 low-transparent latches at the PU's output ports to lessen the effects of these faults. As can be seen in the third column of Table 1, this straightforward action reduces the single-core architecture's overall power usage by 6.7%.

Architecture for Multicore Processors

Fig. 1(b) depicts the multi-core processor design, which consists of N (i.e., 8) PUs with separate IMs. To provide full access to the whole memory area for each PU, a central crossbar link allows each PU to access the 16 shared MBs [14]. This architecture differs from the one that is suggested, in which a proportionately faster cache is shared by multiple slower cores, necessitating a larger supply voltage. Our suggested architecture simplifies the clock-network design and does not call for level-shifters between the cores and the shared cache or an additional faster clock, in contrast to their single, which depends on a fully shared memory-block arrangement. Additionally, the ability to function with a single supply voltage simplifies the design of the entire system and may lead to additional energy savings by avoiding the need for many weakly loaded DC/DC converters. The occasional access conflicts that arise when two or more PUs attempt to access the same MB on the same port are a downside of our methodology.

In this instance, the PU priorities dictate which conflicting requests are fulfilled first, and clock gating is used to halt the waiting PUs in order to prevent needless active power usage. Operating at up to 48 MHz, the multi-core design is likewise optimised for minimum area. All of the cores are in use for processing one lead per core in 761 clock cycles per data for our 8-lead ECG application. When taking into consideration the 8x parallel processing, this translates to a temporal penalty of 12% because of stall cycles as opposed to the number of cycles needed for a single lead in the single-core design. We always modify the multi-core design's clock frequency to match the single-core reference architecture's sampling frequency in order to make up for this penalty when analysing the two architectures' power consumption.

Specifically, we present results for a frequency of 2.3 MHz at nominal supply voltage of 1.2 V. Table 1 illustrates that the clock tree under the suggested architecture only accounts for 5.0% of the total power usage. which, in the end, translates into throughput equivalent to the 16 MHz single-core design. The two columns on the right side of Table 1 offer the corresponding power consumption data. The findings demonstrate that the crossbar overhead, which accounts for just 13% of the total power consumption of the multi-core system, is negligible. By using the same method for glitch mitigation as in the single-core design, this overhead can be further decreased.

Latch placement in the PUs significantly reduces the crossbar interconnect's power consumption, which contributes to the 8.3% power improvement in overall power consumption displayed in Table 1's rightmost column.

Table 1. Distribution of power of the multi-core and single-core designs at 1.2 V power voltage and 16 MHz and 2.3 MHz operating frequencies, respectively, using and without latches in the PU

	Single core		Multi core	
	w / o latches	with latches	w / o latches	with latches
Clock Tree	0.19 mW	0.17 mW	0.24 mW	0.22 mW
MBs	0.24 mW	0.24 mW	0.24 mW	0.24 mW
PUS	2.81 mW	2.81 mW	2.53 mW	2.53 mW
Reduction	-	8.30%	-	6.70%
SL - ICSB	0.48 mW	0.19 mW	0.55 mW	0.33 mW
Total	3.72 mW	3.41 mW	3.56 mW	3.32 mW

Table 2 shows the occupied silicon area for the single- and multi-core designs. As anticipated, the multi-core design's overall PU area is nearly eight times larger than the single-core design's PU area. However, because shared memory accounts for the majority of this area, the overall size of the multi-core architecture is only 1.76 times that of the single-core design.

Table 2. The single-core and multi-core architectures' area results

	Single - core	Multi - core
ICSB	-	20.0 KGE
MBS	576.7 kGE	576.7KGE
PUS	68	kGE541.4 kGE
Topmodule	644.7 kGE	1138.1 kGE

RESULTS

The following is the configuration that we employed for the experiments: The DM of the single-core and multi-core systems pre-stores 1024 samples of an 8-lead ECG signal. The pre-stored ECG samples can be stored for a total of 16 kBytes since each sample takes up 16 bits of memory. In a multi-core design, each core handles one lead, whereas in a single-core design, the leads are processed one after the other. Each lead's results are kept separately in the data memory, and each design requires a total of 16 kBytes of memory to store the results. To investigate the power/performance trade-offs between the two architectures, we run our reference application on each for varying workload needs.

In our experiments, a workload requirement is a number of operations per second (Ops/s). This investigation enables us to extrapolate the findings and patterns to other applications in addition to analysing the architectures for our reference application. Furthermore, we examine the architectures in relation to the ECG sampling frequencies that meet the requirements of our application.

Our operating voltage scaling is restricted to the transistor-threshold level (0.5 V) in order to prevent concerns with performance fluctuation and functional failure, which primarily arise in sub-threshold regions. The processing capacities of the two methods in relation to the supply voltage are displayed in Fig. 2. The single- and multi-core techniques reach 50.1 MOps/s and 343 MOps/s, respectively, at the nominal voltage of 1.2 V. These processing capacities decrease with voltage scaling, as would be predicted. The single-core design can only accomplish up to 806.3 kOps/s when the supply voltage of the designs approaches the threshold level, while the multi-core design can still accomplish up to 5.58 MOps/s.

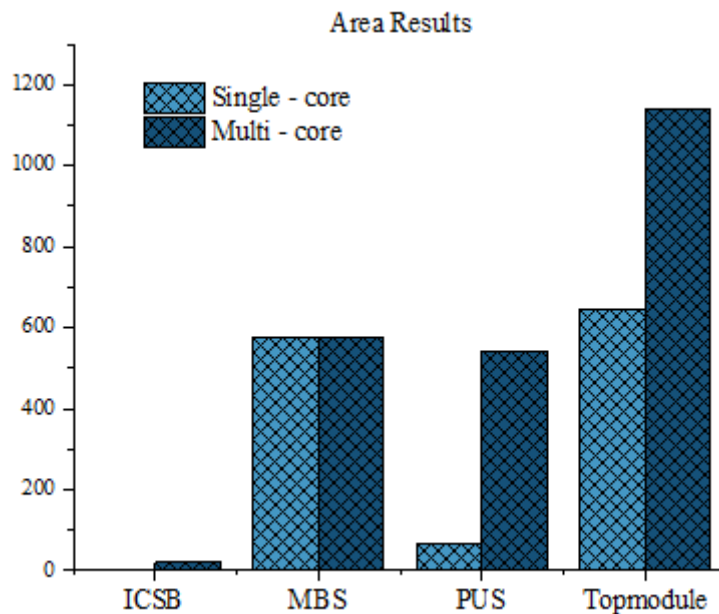


Fig 2. Area results for single core and multi core architectures

The overall power consumption of the single- and multi-core de-sign for varied workload requirements is displayed in Fig. 3(a). The figure illustrates that the multi-core technique is the only workable option for workloads ranging from 50.1 MOps/s to 343 MOps/s. Furthermore, the multi-core design is more energy-efficient than the single-core design when the workload need is between 1356.5 kOps/s and 50.1 MOps/s because it can meet the amount of work needed at a lower operating voltage. Specifically, the single-core design runs at 1.2 V and uses 10.4 mW to achieve a high workload need (50.1 MOps/s), while the multi-core design runs at 0.7 V and uses just 3.5 mW.

As a result, the multi-core approach uses roughly 66% less energy than the single-core architecture. In contrast, the single-core design uses less power than the multi-core design when the required workload is light. This is because the multi-core design can reach the threshold voltage at 5.58 MOps/s workload, while the single-core design can only reach the threshold level at 806.3 kOps/s. More specifically, both designs run at 0.5 V and consume 25.9 μ W for single-core and 28.6 μ W for multi-core in order to achieve the low workload demand of 681 kOps/s. Therefore, compared to the single-core design, the multi-core architecture uses 10.4% more electricity.

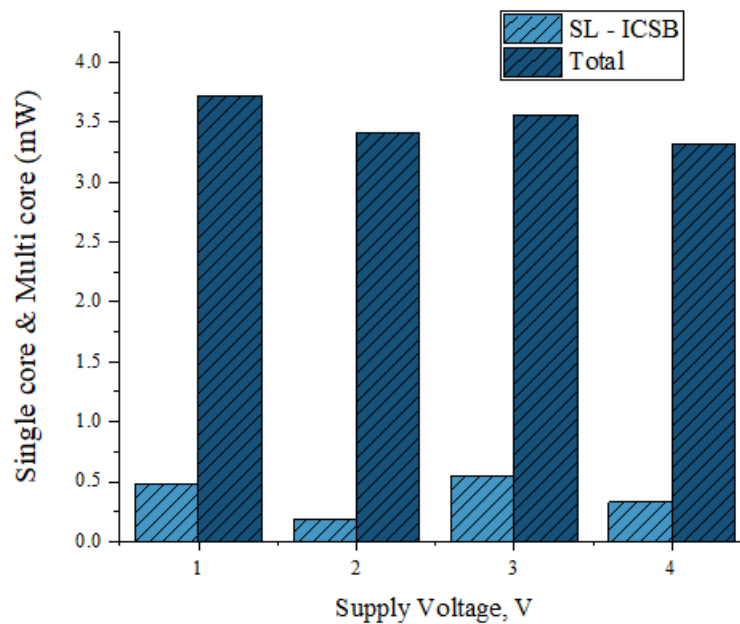


Fig. 3. Maximum permitted ECG sampling rate and matching number of operations for different supply voltages for single-core and multi-core architectures

Our application's associated workload spans 681 kOps/s to 5448 kOps/s, with an ECG sampling rate (fs) ranging from 125 Hz to 1 kHz. The power effectiveness of the multi-core design in comparison to the single-core design for our application is displayed in Fig. 3. The multi-core gets increasingly energy-efficient as the sample rate rises. 55% more power is used by the multi-core design at the maximum sampling rate of fs=1 kHz. However, the multi-core design loses power efficiency if the sampling rate is lowered to 250 Hz. The multi-core uses 10.4% more power than the single core design at fs=125 Hz, the lowest ECG sampling rate in our range.

A fascinating aspect to consider is the contrast between the two designs' dynamic and leaky power consumptions. The single- and multi-core designs' leakage power consumption in our case study is 2.6 μ W and 5 μ W, respectively, with the lowest workload requirement of 681 kOps/s (fs=125 Hz). For the single-core and multi-core systems, respectively, the leakage power consumption amounts to 10% and 17% of the total power consumption.

For different workload demands for the multi-core and single-core architectures, Figs. 4 display the leakage power and dynamic usage of the PCs and the memory, including both IM and MBs. MemDyn denotes the dynamic energy use of the memories in both figures, whereas MemLeak denotes leakage. The PCs' leakage and dynamic power usage are denoted, respectively, by PCsLeak and PCsDyn.




Fig. 4. The multi-core design's power efficiency in comparison to the single-core design for the ECG signal conditioning application

The figures illustrate that at 200 kOps/s and 410 kOps/s, respectively, for the single-core and multi-core systems, MemDyn and MemLeak become equivalent. MemLeak in the multi-core design gets equivalent with the MemDyn power earlier, as expected, due to the multi-core architecture's larger total memory leakage power. Additionally, for the single-core architecture, the total leakage and dynamic power usage become equivalent at about 80 kOps/s, whereas for the multi-core architecture, they are around 140 kOps/s.

5 Conclusion

Due to the comparatively simple and highly efficient computations required for embedded processing of biomedical signals on WBSNs using low-rate physiological information, low voltage actions and parallel processing are made possible. For the processing of biomedical signals on WBSNs, where energy efficiency and immediate processing are critical design requirements, we offer a single- and multi-core processor layout in this study. Using an 8-lead ECG signal conditioning application, we investigated the power/performance trade-offs between the two structures, integrating near threshold voltage computing, for various workloads in order to meet the energy efficiency and data throughput requirements. According to our findings, at high biosignal processing workloads, the multi-core method uses 66% fewer watts than the single-core approach. On the other hand, the multi-core architecture uses 10.4% more power when working on comparatively lesser workloads.

Abbreviation

IM	- Instruction Memory
DM	- Data Memory
PC	- Processor Core

WBSN	- Wireless Body Network
ECG	- Electrocardiogram
NTC	- Near Threshold Computing
PU	- Processing Unit
RISC	- Reduced Instruction Set Computer
SL	- Selection Logic

Competing interests

The authors declare that they have no competing interests.

Consent for publication

Not applicable

Ethics approval and consent to participate

Not applicable

Funding

This research study is sponsored by the institution name. Thank you to this college for supporting this article!

Availability of data and materials

Not applicable

Authors' contribution

Author A supports to find materials and results part in this manuscript. Author B helps to develop literature part.

Acknowledgement

I offer up our fervent prayers to the omnipotent God. I want to express my sincere gratitude to my co-workers for supporting me through all of our challenges and victories to get this task done. I want to express my gratitude for our family's love and support, as well as for their encouragement. Finally, I would like to extend our sincere gratitude to everyone who has assisted us in writing this article.

References

- Aghazadeh, Roghayeh, Javad Frounchi, Fabio Montagna, and Simone Benatti. 2020. “Scalable and Energy Efficient Seizure Detection Based on Direct Use of Compressively-Sensed EEG Data on an Ultra Low Power Multi-Core Architecture.” *Computers in Biology and Medicine* 125 (October): 104004.
- Alsharif, Mohammed H., Abu Jahid, Anabi Hilary Kelechi, and Raju Kannadasan. 2023. “Green IoT: A Review and Future Research Directions.” *Symmetry* 15 (3): 757.
- Bui, Ngoc Thang, Duc Tri Phan, Thanh Phuoc Nguyen, Giang Hoang, Jaeyeop Choi, Quoc Cuong Bui, and Junghwan Oh. 2020. “Real-Time Filtering and ECG Signal Processing Based on Dual-Core Digital Signal Controller System.” *IEEE Sensors Journal* 20 (12): 6492–6503.
- De Giovanni, Elisabetta. 2021. “System-Level Design of Adaptive Wearable Sensors for Health and Wellness Monitoring.” Lausanne, EPFL. <https://doi.org/10.5075/EPFL-THESIS-8052>.
- De Giovanni, Elisabetta, Farnaz Forooghifar, Gregoire Surrel, Tomas Teijeiro, Miguel Peon, Amir Aminifar, and David Atienza Alonso. 2023. “Intelligent Edge Biomedical Sensors in the Internet of Things (IoT) Era.” In *Emerging Computing: From Devices to Systems: Looking Beyond Moore and Von Neumann*, edited by Mohamed M. Sabry Aly and Anupam Chattopadhyay, 407–33. Singapore: Springer Nature Singapore.
- De Giovanni, Elisabetta, Fabio Montagna, Benoît W. Denking, Simone Machetti, Miguel Peón-Quirós, Simone Benatti, Davide Rossi, Luca Benini, and David Atienza. 2020. “Modular Design and Optimization of Biomedical Applications for Ultralow Power Heterogeneous Platforms.” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 39 (11): 3821–32.
- De Giovanni, Elisabetta, Adriana Arza Valdés, Miguel Peón-Quirós, Amir Aminifar, and David Atienza. 01 Oct.-Dec 2021. “Real-Time Personalized Atrial Fibrillation Prediction on Multi-Core Wearable Sensors.” *IEEE Transactions on Emerging Topics in Computing* 9 (4): 1654–66.
- Denking, Benoît Walter. 2023. “Exploring Brain-Inspired Multi-Core Heterogeneous Hardware Templates for Low-Power Biomedical Embedded Systems.” Lausanne, EPFL. <https://doi.org/10.5075/EPFL-THESIS-9353>.
- Djelouat, Hamza, Mohamed Al Disi, Issam Boukhenoufa, Abbes Amira, Faycal Bensaali, Christos Kotronis, Elena Politi, Mara Nikolaidou, and George Dimitrakopoulos. 2020. “Real-Time ECG Monitoring Using Compressive Sensing on a Heterogeneous Multicore Edge-Device.” *Microprocessors and Microsystems* 72 (February): 102839.
- Ingolfsson, Thorir Mar, Andrea Cossetini, Xiaying Wang, Enrico Tabanelli, Giuseppe Tagliavini, Philippe Ryvlin, Luca Benini, and Simone Benatti. 2021. “Towards Long-Term Non-Invasive Monitoring for Epilepsy via Wearable EEG Devices.” In *2021 IEEE Biomedical Circuits and Systems Conference (BioCAS)*, 01–04. IEEE.
- Kartsch, Victor, Marco Guermandi, Simone Benatti, Fabio Montagna, and Luca Benini. 2019. “An Energy-Efficient IoT Node for HMI Applications Based on an Ultra-Low Power Multicore Processor.” In *2019 IEEE Sensors Applications Symposium (SAS)*, 1–6. IEEE.

- Kartsch, Victor, Giuseppe Tagliavini, Marco Guermandi, Simone Benatti, Davide Rossi, and Luca Benini. 2019. "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* 13 (5): 893–906.
- Mane, Shreya. 2023. "Theoretical Study on Embedded Processor and Networking." *International Journal of Engineering Technology and Management Sciences* 7 (3): 861–67.
- Paulin, Gianna, Renzo Andri, Francesco Conti, and Luca Benini. 2021. "RNN-Based Radio Resource Management on Multicore RISC-V Accelerator Architectures." *IEEE Transactions on Very Large Scale Integration Systems* 29 (9): 1624–37.
- Prasad, Rohit, Satyajit Das, Kevin J. M. Martin, and Philippe Coussy. 2021. "Floating Point CGRA Based Ultra-Low Power DSP Accelerator." *Journal of Signal Processing Systems* 93 (10): 1159–71.
- Rossi, Davide, Francesco Conti, Manuel Eggiman, Alfio Di Mauro, Giuseppe Tagliavini, Stefan Mach, Marco Guermandi, et al. 2022. "Vega: A Ten-Core SoC for IoT Endnodes With DNN Acceleration and Cognitive Wake-Up From MRAM-Based State-Retentive Sleep Mode." *IEEE Journal of Solid-State Circuits* 57 (1): 127–39.

- Sharifshazileh, Mohammadali, Karla Burelo, Johannes Sarnthein, and Giacomo Indiveri. 2021. “An Electronic Neuromorphic System for Real-Time Detection of High Frequency Oscillations (HFO) in Intracranial EEG.” *Nature Communications* 12 (1): 3095.
- Wang, Xiaying, Michele Magno, Lukas Cavigelli, and Luca Benini. 2020. “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* 7 (5): 4403–17.