

A Dual Processor Energy-Efficient Platform with Multi-core Accelerator for Smart Sensing

Antonio Pullini^(✉), Stefan Mach, Michele Magno, and Luca Benini

Integrated Systems Laboratory, ETH Zurich, Gloriastrasse, 35, 8092 Zurich, Switzerland
{antonio.pullini, stefan.mach, michele.magno,
luca.benini}@iis.ee.ethz.ch

Abstract. Energy-efficient computing has increasingly come into focus of research and industry over the last decade. Ultra-low-power architectures are a requirement for distributed sensing, wearable electronics, Internet of Things and consumer electronics. In this paper, we present a dual-mode platform that includes an ultra-low power Cortex Arm M4 microcontroller coupled with a highly energy efficient multi-core parallel processor. The platform is designed to maximize the energy efficiency in sensors applications by exploiting the Cortex Arm M4 to achieve ultra-low power processing and power management, and enables the multi-core processor to provide additional computational power for near-sensor data centric processing (i.e. accelerating Convolutional Neural Networks for image classification) increasing energy efficiency. The proposed platform enhances the application scenarios where on-board processing (i.e. without streaming out the sensor data) enables intensive computation to extract complex features. The platform is geared towards applications with limited energy budget, as for example in mobile or wearable scenarios where the devices are supplied by a battery. Experimental results confirm the energy efficiency of the platform, demonstrate the low power consumption, and the benefits of combining the two processing engines. Compared to a pure microcontroller platform we provide a boost of 80× in terms of computational power when running general purpose code and a boost of 560× when performing convolutions. Within a reasonable power budget of 20 mW compatible to battery-operated scenarios the system can perform 345 MOPS of general purpose code or 1.5 GOPS of convolutions.

Keywords: Low power design · Sensors platform · Energy efficiency · Power management · Multi-core processor

1 Introduction

Due to the vast improvements in sensors technology, digital processors, device miniaturization and thanks to the availability of ubiquitous wireless connectivity, intelligent sensor devices are becoming increasingly smart and this leads to always-on connected products. The Internet of Things (IoT) paradigm, promises to have trillions of those sensors devices in nearly future deployed around the world [1]. Partially this revolution has already started, and today sensors devices are gaining immense popularity, with

people increasingly surrounded by “smart” objects, from phones to clothing, from glasses to watches, finding applications from home automation to healthcare [2].

The IoT is also creating formidable research challenges. In particular, trillions of sensors will produce huge amounts of data that need to be sent and stored somewhere. Moreover, the data by themselves do not provide value unless we can turn them into actionable, contextualized information. In fact, to produce useful information the data needs to be processed by some intelligent system somewhere along the line. Big-data mining techniques allow us to gain new insights by batch-processing and off-line analysis. Machine learning technologies are used with great success in many application areas, solving real-world problems in entertainment systems, robotics, health care, and surveillance [1]. More and more researchers are tackling classification and decision-making problems with the help of brain-inspired algorithms, featuring *many stages* of feature extractors and classifiers with lots of parameters that are optimized using the unprecedented wealth of training data that is currently available. However, machine learning requires complex software and significant computational power to be really effective [4, 5].

Today there are many IoT applications that use a centralized approach where data processing is done far from the sensor. In these applications, data is sent directly to a remote host capable of running complex and power hungry algorithms. This is, for example, the approach used by the cloud computing adopted by the biggest service companies as Google and Facebook and millions of users [6]. It is clear that, as the number of data generating remote sensors increase steadily, the communication infrastructure will be not sufficient to deal with the enormous amounts of data being generated all over the planet. For a truly scalable and robust IoT infrastructure to succeed, in-situ, close to the sensor and distributed real-time feature extraction, analysis, classification, and local decision-making are essential [7].

In recent years, there have been many research efforts to design new processors to match the requirements of computational resources required by in-situ signal processing with low power consumption needed for operating long-lasting sensors devices [7–10]. There are two approaches to improve the performance of ultra-low-power processors that have shown promise. The first one is to exploit parallelism as much as possible. Parallel architectures for near-threshold operation, based on multi-core clusters, have been explored in recent years with different application workloads for an implementation in a 90 nm technology [17]. A second very prolific research area is exploiting low-power fixed-function hardware accelerators coupled with programmable parallel processors to retain flexibility while improving energy efficiency for specific workloads [11, 12]. Such near-threshold parallel heterogeneous computing approaches hold great promise.

In this work, we present a complete hardware platform that includes an heterogeneous multi-core System on Chip (SoC), capable of operating on a wide voltage range, paired with an ultra-low power ARM Cortex M4 micro controller that are able to interface to a wide set of sensors. The platform is designed to achieve the best energy efficiency for a wide range of applications by combining the ultra-low power of the highly integrated ARM microcontroller and the powerful multicore SoC. The ARM M4, which is designed for battery powered applications such as wearable electronics, is used to configure the SoC processor, provides the power management of the board and can also

process preliminary sensor data. In this way, the SoC processor is activated by the ARM M4 only when it has not enough computational resources (i.e. processing convolutional neural network for video processing) or in cases when it is more energy efficient to process the data on the SoC (i.e. if the SoC can accelerate the algorithm by a factor of 10× or more). The platform has been designed as a generic testbed and supports several peripherals where sensors can be attached.

The rest of the paper is organized as follows: Sect. 2 presents related work, Sect. 3 describes the SoC ultra low-power multi-core parallel platform (PULP), Sect. 4 illustrates the multi-processor platform that has been designed and developed, Sect. 5 shows the experimental results and Sect. 6 concludes the paper.

2 Related Works

Research on intelligent sensors systems has been very prolific in recent years with a variety of solutions in a wide range of application scenarios [1, 3]. There are many examples of implemented and deployed wearable devices that attempt to exploit intelligent sensing and wireless communication to monitor human activities [13–15]. The main challenges of IoT devices design are to prolong the operating lifetime and to enhance usability, maintenance, and mobility, while keeping a small and unobtrusive form factor [2]. Many IoT devices such as for example mobile and wearable sensing systems have to provide continuous data monitoring, acquisition, processing, and classification. Supporting such continuous operation using only ultra-small batteries poses unique challenges in energy efficiency [16].

As sensor data processing based on machine learning needs computational performance, most IoT applications today on the market have focused on using smartphones as a centralized hub that provides a powerful computing platform and allows a network of smaller sensors to be connected. Pushing on energy efficiency, state-of-the-art commercial ultra-low power processors are trying to exploit novel solutions to extract as much as possible out of silicon. A novel approach to further improve the energy efficiency is near-threshold computing, which exploits the fact that CMOS technology is most efficient when operated near the device voltage threshold [19]. In particular in [20–22] the authors show examples of near-threshold ultra-low power microcontrollers, with the latter also exploiting SIMD parallelism to improve performance. There are microcontrollers that can embed custom hardware accelerators [23] as well to improve the computational performance. However, such approaches limit the flexibility of the solution affecting the cost and scale economy. In this paper, we show the potential of combining an ARM Cortex M4 with a state-of-the-art multi-core accelerator in a single platform to maximize the energy efficiency of a wide range of sensor applications providing extraordinary computational power.

3 Pulp Overview

The main aim of this work is to build a multi-modal and multi-processor sensing platform that embeds an ultra-low-power parallel-processor called PULP (Parallel Ultra Low

Power Platform) [17]. The PULP processor has been designed specifically to take advantage of the energy-efficient near-threshold regime. The degradation of performance, caused by the aggressive voltage scaling, is recovered by increasing the parallelism. The PULP platform is built upon a cluster of tightly coupled cores. To avoid the huge overhead of a cache coherent system, the cores do not have private data caches but share data through a Tightly Coupled Data Memory (TCDM). The TCDM is composed of several single ported memory cuts connected to the processors with a non-blocking logarithmic interconnect. The interconnect grants single cycle access when there is no contention and, by using appropriate banking factors and interleaving, on average the access contention remains below 10% even for load/store intensive applications. The Instruction-Cache (I-Cache) is shared among all cores and is implemented with Standard Cell Memory (SCM) cuts to optimize the energy of instruction fetching. Data transfer between L1 TCDM and the main SoC memory is done by a system DMA capable of queuing multiple transfers with ultra-low latency programming interface dedicated to each core. The system is completely event based: the cores, when waiting for synchronization events or for I/O, are forced in an idle state by a dedicated hardware Event Unit. The event unit performs the gating of the cores and provides hardware support for fast core synchronization (Fig. 1).

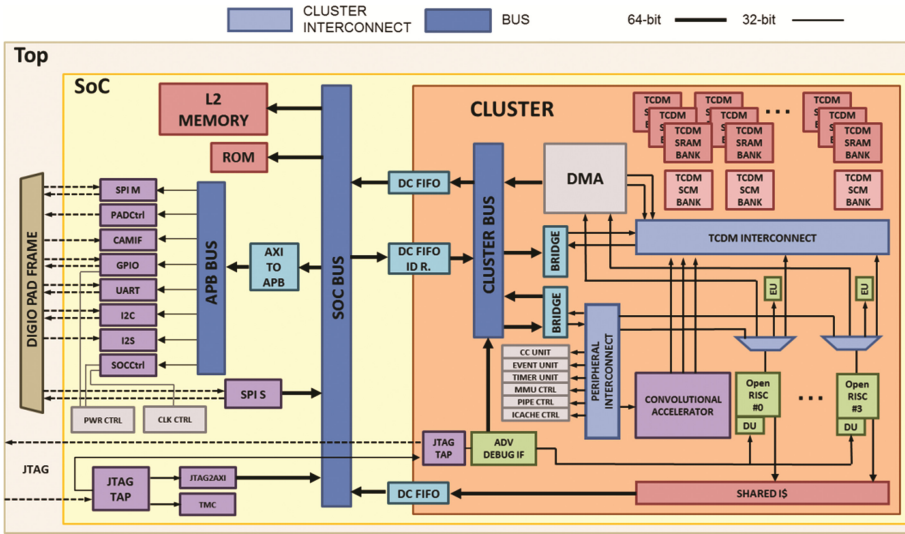


Fig. 1. MiaWallace architectural diagram

The SoC named MiaWallace in this work is an implementation of the PULP platform in UMC 65 nm with the addition of a convolutional hardware accelerator. It has four cores and features an L1 TCDM of 80 KB (64 KB SRAM and 16 KB SCM based), a 4 KB instruction cached based entirely on SCM, a 256 KB L2 memory and a full set of peripherals. The cores are compliant with the OpenRISC ISA, with instruction set extensions for DSP applications to improve performance [24].

The dedicated Hardware Convolution Engine (HWCE) is directly connected to the L1 TCDM memory through the logarithmic interconnect just like processor cores. It uses three dedicated ports toward the shared memory to sustain the full bandwidth required by its engine. Its core is made of two sum of products unit, each of which is capable of performing a 5×5 convolution on 16 bits input data. Although it is optimized for 5×5 convolutions it can, with a little loss of performance perform convolutions of different sizes allowing applications that use convolutions (such as convolutional neural networks) to be processed efficiently [12].

The PULP includes two SPI (Serial Peripheral Interface) interfaces (one master and one slave), I2C, I2S, GPIOs, a boot ROM and a JTAG interface suitable for testing purposes. Both SPI interfaces can be configured in *single* mode or *quad* mode depending on the required bandwidth, and they are suitable for interfacing the SoC with a large set of off-chip components (non-volatile memories, voltage regulators, cameras, etc.).

PULP is able to operate in two different modes: *slave* mode or *stand-alone* mode. When configured in slave mode, PULP behaves as a many-core accelerator of a standard *host* processor (e.g. an ARM Cortex M4 low-power microcontroller). In this configuration the host microcontroller is responsible for loading the application and the data on to the PULP L2 memory through the SPI master interface. After this the microcontroller initiates and synchronizes the computation through dedicated memory mapped signals (e.g. fetch enable) and GPIOs. When configured in stand-alone mode, the boot code in the on-chip ROM is able to detect a flash memory on its SPI master interface and, if present, will load the program to the L2 memory and starts the execution.

The SoC is divided in two voltage domains, one for the cluster and one for the peripherals and L2 memories. The cluster works on a wide range of voltages starting as low as 0.62 V while the minimum operating voltage of the peripheral domain is limited by the L2, whose performance degrades severely below 0.8 V. Figure 2 shows the chip micrograph and a table with the main features.

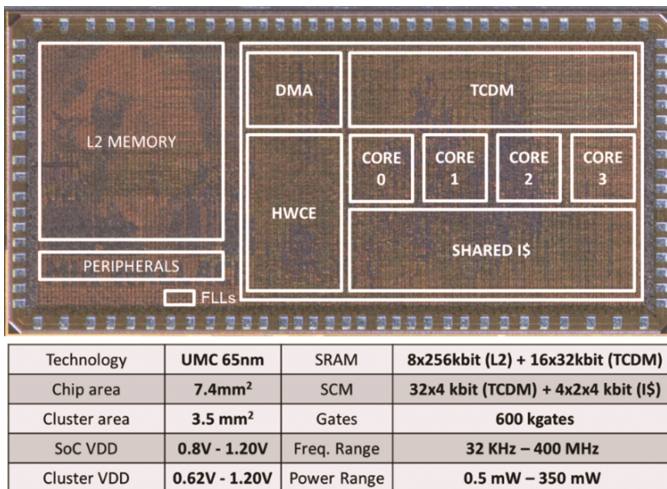


Fig. 2. Chip micrograph and main features

4 Platform Architecture

Figure 3 shows the architectural block diagram of the implemented platform. The designed multi-processor platform features an ARM Cortex M4 ultra low-power microcontroller as well as the Mia Wallace SoC as the main architectural blocks. The two devices are interconnected by various interfaces. In particular, PULP’s slave SPI interface is driven by the microcontroller to have access to the entire memory space of the PULP system. With this interface, the microcontroller can assume the role of a host controller using PULP as an accelerator. Additionally, a shared I²C Bus and GPIO connections between the two devices allow for user-programmable signaling or data exchange.

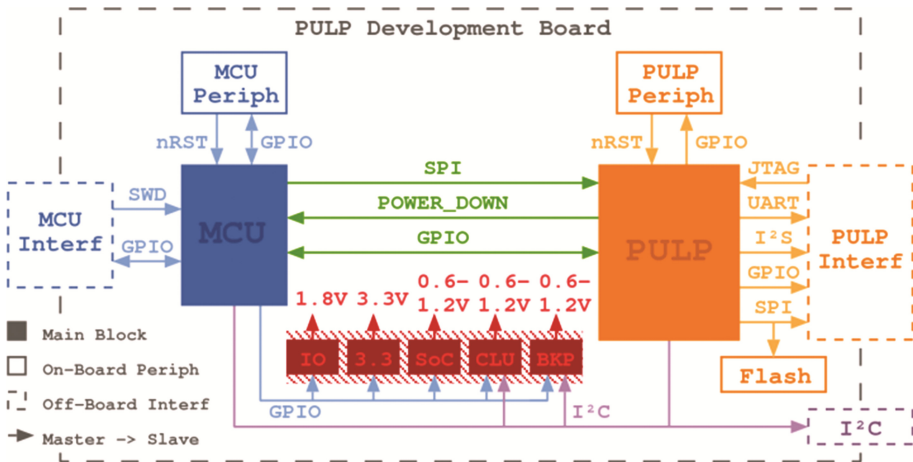


Fig. 3. Sytem architecture

Both PULP and the microcontroller are connected to a set of LEDs and push buttons to allow for basic user interaction. A flash memory for loading PULP programs in stand-alone mode is part of the platform. Although in the current version of the platform we didn’t embed any sensors, all interfaces of both devices are accessible via pin headers and connectors so that a multitude of sensors or other peripheral devices can be attached to the system. Thus, serial interfaces such as I²C, UART, SPI, QUAD-SPI, I²S and GPIOs are all available for sensors board extension.

As the platform is mainly targeting mobile and wearable applications, to ensure simplicity and portability, the platform is powered from a single power source such as a laboratory power supply or a single Li-Po battery cell. A set of DC-DC converters are a part of the platform in order to provide the necessary supply voltages for PULP, the microcontroller and peripherals to be attached to the system.

All the supplies can be controlled by both the microcontroller and the MiaWallace SoC, however to achieve ultra-low power consumption the microcontroller can completely shut down MiaWallace and ensure correct wakeup after deep sleep modes.

To improve the energy efficiency of the platform and provide flexibility, different power supplies have been optimized for different current ranges and average on times.

Table 1 shows the main features of the power converters we have evaluated. High Efficiency Power Converter (HEPCO) has been designed to extend the voltage range of commercially available DC/DC components while keeping the maximum possible efficiency.

Table 1. Efficiency and Ripple for different DC/DC converters

Converter	1.2 V				0.6 V			
	0.1 mA		10 mA		0.1 mA		10 mA	
	Eff (%)	Rip (V)	Eff	Rip	Eff	Rip	Eff	Rip
TPS62080 PFM/PWM	40	0.4	80	0.2			73	0.2
TPS62080snooze mode	57	4.2	79	0.9	48	4.8	71	0.8
TPS62361B	30	0.9	84	0.5	23	2.4	78	0.8
HEPCO w/TPS62736	82	1.9	88	0.5	74	3.9	80	1.7
HEPCO w/TPS62737	76	3.2	79	2.9	63	9.2	68	8.0

For the cluster domain we choose the TPS62361B, which has an operating range compatible with our requirements and it has a very fine tuning range to find the optimal operating point for a given application. The minimum current of the cluster is above the point for which this converter starts to lose efficiency. For the peripheral domain of the MiaWallace we use the HEPCO implemented with the TPS62736 due to its high efficiency at lower currents.

The choice of the microcontroller was driven by the low power features available and the flexibility on the power modes. After comparing various microcontrollers available today on the market we choose the Ambiq Apollo MCU. The microcontroller combines ultra-low-power sensor conversion electronics with a 32-bit ARM Cortex-M4F processor. It also integrates 512 KB of flash memory, 64 KB of RAM and a Floating Point Unit which is a big advantage compared to the other MCUs in the ultra-low power world.

Other main components of the Apollo MCU include: 10 bit ADC with 8 channels, temperature sensor, I2C/SPI interface, 50 GPIO, and one UART. Furthermore, the Apollo MCU includes a set of timing peripherals based on Ambiq's AM08XX and AM18XX Real-Time Clock (RTC) families. The RTC, timers, and counters may be driven by three different clock sources: a low frequency RC oscillator, a high frequency RC oscillator, and a 32.768 kHz crystal (XTAL) oscillator. With its extremely low active mode power of $<40 \mu\text{A}/\text{MHz}$, it is possible to perform complex sensor processing algorithms on the Apollo MCU. The Apollo MCU also includes a Power Management Unit (PMU) that controls the transitions of the MCU between the following power modes:

Active mode: in this state, the processor M4F is switched on, all clocks are active and instructions are being executed. The MCU will return to active mode during reset, when an interrupt is received by the **Nested Vectored Interrupt Controller** or a Debug Event is received.

Sleep mode: during this mode, the M4F is powered up, the clocks (HCLK, FCLK) are not active. The difference between this state and the Deep Sleep Mode is that the M4F logic is still on and it can return to Active State rapidly on a wakeup event.

Deep Sleep mode: in this state, the M4F enters a State Retention Power Gating (SRPG) where the main power is removed, but the registers in the MCU retain their values. The clocks are not active, and the clock sources for HCLK and FCLK can be deactivated. Table 2 shows measurements performed on the MCU during different operating modes relevant for the project.

Table 2. Ambiq Apollo operating modes and current consumption

Scenario	Power consumption
Deep sleep, RTC disabled	100 nA @ 2 V
Deep sleep, 8/64 KB ram block retention, RTC on 1 s, incrementing one variable	125 nA @ 2 V
Deep sleep, 64/64 KB ram block retention, RTC on 1 s, incrementing one variable	435 nA @ 2 V
Normal sleep, RTC disabled, 8/64 KB ram block retention	50 μ A @ 2 V
Active mode, 64/64 KB ram retention	1.3 mA @ 2 V

5 Experimental Results

5.1 Experimental Setup

The platform has been designed and implemented on a small-outline PCB just 8.7 cm \times 5.7 cm in size and is shown in Fig. 4. The whole platform can then be supplied by a single battery.




Fig. 4. PCB photo

Before measurements take place, programs for PULP and the microcontroller are loaded into PULP's flash memory and the ARM's onboard flash via the JTAG and SWD debug ports, respectively. The platform is supplied by and measurements are then taken

with the Keysight N6705B DC power analyzer. This approach allows for precise measurements of dissipated power in the individual components of the platform as well as for the calculation of converter efficiencies.

5.2 Computational Performance

MiaWallace with its wide operating range and the availability of the HWCE enables multiple working modes that could cover many applications typical of the IoT domain. Detailed measurements of performances and power consumption of the SoC in stand-alone mode have been performed using an Advantest SoCV93000 ASIC tester. Figure 5a shows the maximum operating frequency for different voltage points over the whole operating range, while Fig. 5b shows how the energy efficiency changes in the same range. When the HWCE is on, we consider the power in different conditions depending on the computation over communication ratio (CCR) which depends highly on the topology of the CNN. The number of operation per second is assuming that the cores can perform 1 instruction per cycle. This assumption is true for many DSP kernels especially where the operands are smaller than 32 bits and the system can benefit from the available vector support extensions.

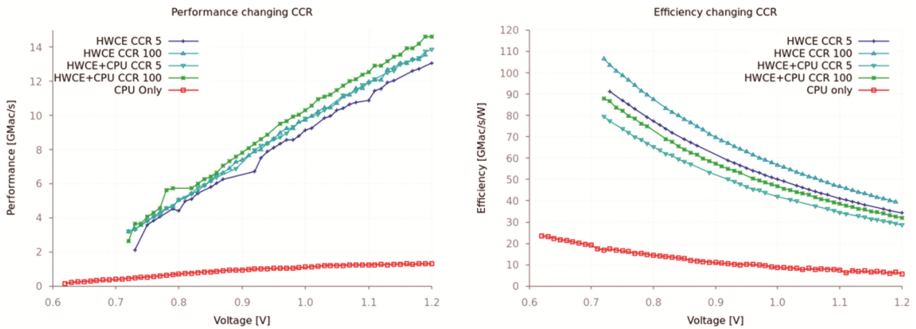


Fig. 5. Performance and efficiency of the SoC for different operating points

Both figures clearly show the boost in efficiency and performance given by the accelerator. A high accuracy CNN architecture as GoogleLeNet requires nearly 2.5 GOPs to process a frame of 320×240 pixels. The peak performance curve shows that even at modest voltages the system can sustain a full blown convolutional neural network when using the HWCE. From both graphs we can see that when using the CPU only the system can operate at very low voltages. This is possible thanks to the heterogeneity of the L1 TCDM and the use of the SCMs. This ULP mode is very useful for example in all the applications where the environment has to be monitored continuously (light, noise, temperature) and only upon an event the full processing is performed. The system can sustain a maximum of 14 GMAC/s at 1.2 V and reach an energy efficiency of 108 GMAC/s/W.

5.3 System Profiling

After measuring the SoC and the power supply in isolation we implemented a simple power management firmware in the Apollo MCU and we measured the whole system in different operation points and during a synthetic application in which we moved through the different power modes. Figure 6 shows the profile with the states and the measurements results of the whole sequence while Table 3 gives more details about the states.

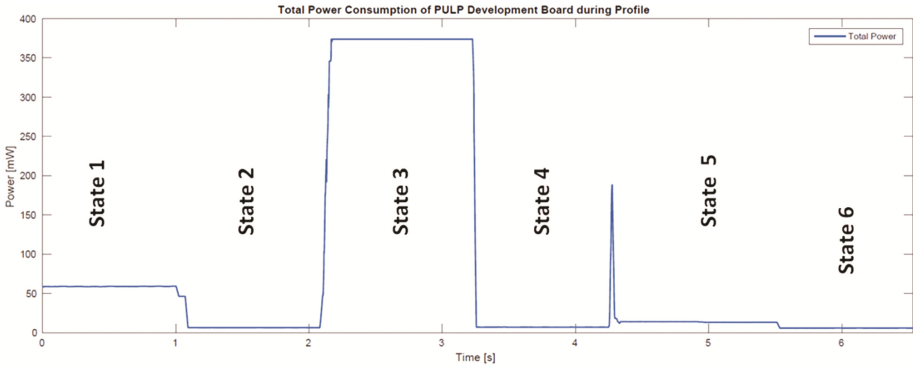


Fig. 6. Power profiling of the whole platform tested in our lab

Table 3. State Details

State	Ambiq Apollo state			MiaWallace SoC state			MiaWallace cluster state			Conv Eff%	Perf GMAC/s
	Vdd V	Freq Mhz	Pwr mW	Vdd V	Freq Mhz	Pwr mW	Vdd V	Freq Mhz	Pwr mW		
1	2.1	20	0.2	1.13	50	9.1	1.16	50	38	85.8	1.8
2	2.1	20	0.3	1.13	50	0.7	OFF	n/a	<1u	91.1	n/a
3	2.1	20	0.3	1.13	200	34.5	1.16	350	330	86	12.5
4	2.1	20	0.3	1.13	200	1.3	OFF	n/a	<1u	87.2	n/a
5	2.1	20	0.2	0.79	5	0.7	0.68	5	1.8	79.3	0.02
6	2.1	20	0.3	1.13	5	0.2	OFF	n/a	<1u	90	n/a

6 Conclusions

In this paper we presented a multi-modal multi-processor platform designed to maximize the energy efficiency of smart sensing applications. The platform can be supplied by a single battery and can host a wide range of sensors through several peripherals. The platform exploits the combination of the two processors to achieve energy efficiency. In particular, with the ultra-low power commercial microcontroller, it is possible to achieve very low power states and manage the power supply of the rest of the platform.

On the other hand, the presence of the multi-core energy efficient accelerator brings extraordinary computational resources even when working in a very tight power envelope. The platform has been designed carefully also to achieve conversion efficiency on the power domains needed for the PULP processor. Experimental results on the developed platform shows the energy efficiency and the low power of the platform. The platform is ready to host sensors and applications that will be studied in future works.

References

1. Gubbi, J., Buyya, R., Marusic, S., Palaniswami, M.: Internet of Things (IoT): a vision, architectural elements, and future directions. *Future Gener. Comput. Syst.* **29**(7), 1645–1660 (2013)
2. Da Xu, L., He, W., Li, S.: Internet of Things in industries: a survey. *IEEE Trans. Ind. Inform.* **10**(4), 2233–2243 (2014)
3. Govindaraju, V., Rao, C.: *Machine Learning: Theory and Applications*. Elsevier, Amsterdam (2013)
4. Michalski, R.S., Carbonell, J.G., Mitchell, T.M. (eds.) *Machine learning: an artificial convolutional networks*. In: *Proceedings of the 52nd Annual Design Automation Conference*, p. 108. ACM (2013)
5. Hashem, I.A.T., Yaqoob, I., Anuar, N.B., Mokhtar, S., Gani, A., Khan, S.U.: The rise of “big data” on cloud computing: review and open research issues. *Inf. Syst.* **47**, 98–115 (2015)
6. Hwang, K., Dongarra, J., Fox, G.C.: *Distributed and Cloud Computing: From Parallel Processing to the Internet of Things*. Morgan Kaufmann, Boston (2013)
7. Kahng, A.B., Kang, S., Kumar, R., Sartori, J.: Enhancing the efficiency of energy-constrained DVFS designs. *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.* **21**(10), 1769–1782 (2013)
8. Wang, Z., Liu, Y., Sun, Y., Li, Y., Zhang, D., Yang, H.: An energy-efficient heterogeneous dual-core processor for Internet of Things. In: *2015 IEEE International Symposium on Circuits and Systems (ISCAS)*, Lisbon (2015)
9. Dreslinski, et al.: Centip3De: a 64-core, 3D stacked, near-threshold system. *IEEE Micro* **33**(2), 8–16 (2013)
10. Jeon, D., Kim, Y., Lee, I., Zhang, Z., Blaauw, D., Sylvester, D.: A 470 mV 2.7 mW feature extraction-accelerator for micro-autonomous vehicle navigation in 28 nm CMOS. In: *Proceedings of 2013 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, pp. 166–168 (2013)
11. Yoon, J.-S., Kim, J.-H., Kim, H.-E., Lee, W.-Y., Kim, S.-H., Chung, K., Park, J.-S., Kim, L.-S.: A unified graphics and vision processor with a 0.89 uW/fps pose estimation engine for augmented reality. *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.* **21**(2), 206–216 (2013)
12. Ghasemzadeh, H., Jafari, R.: Ultra low-power signal processing in wearable monitoring systems: a tiered screening architecture with optimal bit resolution. *ACM Trans. Embed. Comput. Syst. (TECS)* **13**(1) (2013). Article No. 9
13. Magno, M., Brunelli, D., Sigrist, L., Andri, R., Cavigelli, L., Gomez, A., Benini, L.: InfiniTime: multi-sensor wearable bracelet with human body harvesting. *Sustain. Comput. Inform. Syst.* **11**, 38–49 (2016)
14. Cavigelli, L., Magno, M., Benini, L.: Accelerating real-time embedded scene labeling with convolutional networks. In: *Proceedings of the 52nd Annual Design Automation Conference*, p. 108. ACM, June 2015
15. Magno, M., Spagnol, C., Benini, L., Popovici, E.: A low power wireless node for contact and contactless heart monitoring. *Microelectron. J.* **45**(12), 1656–1664 (2014)

16. Magno, M., Salvatore, G.A., Mutter, S., Farrukh, W., Troester, G., Benini, L.: Autonomous smartwatch with flexible sensors for accurate and continuous mapping of skin temperature. In: 2016 IEEE International Symposium on Circuits and Systems (ISCAS), pp. 337–340. IEEE, May 2016
17. Rossi, D., et al.: A -1.8 V to 0.9 V body bias, 60 GOPS/W 4-core cluster in low-power 28 nm UTBB FD-SOI technology. In: SOI-3D-Subthreshold Microelectronics Technology Unified Conference (S3S). IEEE, Rohnert Park (2015)
18. Conti, F., Benini, L.: A ultra-low-energy convolution engine for fast brain-inspired vision in multicore clusters. In: 2015 Design, Automation and Test in Europe Conference and Exhibition (DATE), Grenoble, pp. 683–688 (2015)
19. Dreslinski, R., Wiecekowsi, M., Blaauw, D., Sylvester, D., Mudge, T.: Near-threshold computing: reclaiming Moore’s law through energy efficient integrated circuits. In: Proceedings of the IEEE, vol. 98, pp. 253–266, February 2010
20. Ickes, N., Sinangil, Y., Pappalardo, F., Guidetti, E., Chandrakasan, A.P.: A 10 pJ/cycle ultra-low-voltage 32-bit microprocessor system-on-chip. In: 2011 Proceedings of the ESSCIRC (ESSCIRC), pp. 159–162. IEEE, September 2011
21. Bol, D., De Vos, J., Hocquet, C., Botman, F., Durvaux, F., Boyd, S., Flandre, D., Legat, J.-D.: SleepWalker: a 25-MHz 0.4-V Sub-mm² 7-uW/MHz microcontroller in 65-nm LP/GP CMOS for low-carbon wireless sensor nodes. *IEEE J. Solid-State Circ.* **48**, 20–32 (2013)
22. Botman, F., Vos, J.D., Bernard, S., Stas, F., Legat, J.-D., Bol, D.: Bellevue: a 50 MHz variable-width SIMD 32 bit microcontroller at 0.37 V for processing-intensive wireless sensor nodes. In: Proceedings of 2014 IEEE Symposium on Circuits and Systems, pp. 1207–1210 (2014)
23. Fujita, T., Tanaka, T., Sonoda, K., Kanda, K., Maenaka, K.: Ultra low power ASIC for R-R interval extraction on wearable health monitoring system. In: 2013 IEEE International Conference on Systems, Man, and Cybernetics (SMC), pp. 3780–3783, October 2013
24. Gautschi, M., et al.: Tailoring instruction-set extensions for an ultra-low power tightly-coupled cluster of OpenRISC cores. In: 2015 IFIP/IEEE International Conference on Very Large Scale Integration (VLSI-SoC), Daejeon, pp. 25–30 (2015)