Analysis of the Performance of the New Generation of 32-bit Microcontrollers for IoT and Big Data Application

Jovan Ivković*, Jelena Lužija Ivković**
* ITS Visoka škola strukovnih studija za IT/ Informacione tehnologije, Belgrade, Serbia
** MyITLab, Belgrade, Serbia
jovan.eps@gmail.com, jelena.luzija@gmail.com

***Abstract* — This paper explores the relationship between performance of different generations of MCU solutions for classical computing processing, as well as FPU/DSP processing of input information. Finally, it questions how fast software-controlled I/O ports (GPIO) are. By using the battery of standardized computer tests, and custom GPIO, GPU / DSP tests, the report gives clear and measurable results. The results of measuring carried out in the study clearly show that it is possible to unload the network resources necessary for transmitting large quantities of information if they represent a part of a complex pre-treatment data for Big Data processing systems on the MCU platforms themselves with IoT sensor devices.**

***Keywords:* Microcontrollers - MCU, Internet of Things - IoT, SoC development systems, Cyber physical systems, Big Data, Sensor system, Sensor networks, RISC CPU Architecture, Floating point unit – FPU, Digital signal processing DSP, ARM, PIC32, Xtensa LX106, ESP32, NodeMCU IoT development boards, Arduino development boards, GPIO.**

I. INTRODUCTION

In early 2016, IoT in its scope and complexity exceeded the Internet PC - mobile world. According to the latest estimates, the world is now using about ten billion devices and computers, while IoT devices are numbering more than a hundred billion units [1]. If we add diversity by which PCs (excluding servers) have two different forms, while mobile devices are found in dozens, IoT devices exist in millions of forms and shapes and are a leader in this area as well. The exponential increase in the number of IoT / sensor system inevitably leads to the need for higher volume transmission and processing of information. The existing network infrastructure and Internet connections with Big Data systems become a bottleneck of IoT - Big Data architecture. This has led to the need that treatment and processing is executed locally from within the IoT / sensor networks. Intel and other leading manufacturers have embraced this idea under different names: local, on-premises, before/pre-processing and likewise.

Since 2014, with the emergence of a new generation of high-performance 32-bit RISC microcontrollers based on the ARM Cortex-M4-7, MIPS 32 and Tensilica Xtensa LX106 architectures, adding FPU / DSP cores and support for network and wireless technology, have made a great progress in the field of transfer functions with SoC / FPGA system on these solutions. Previous works [2] [3] [4] show, that computer systems based on low power multi-core SoC CPUs with HMP can take role of pre-processing a very fast local sensor network data prior to sending them to Big Data system, with new advance in the area of MCU's performance. The central question in this researc was whether this new generation of MCUs has sufficient processing power to take part of the pre-processing (IoT gateway) role.

II. CYBER PHYSICAL SYSTEMS (CPS) AND INTELIGENT DEVICES

Current CPS concept can be described as integration of sensor networks embedded with microcontrollers and SoC systems with provided internet connection (IoT) on local side to cloud computing infrastructure that process, analyze, manage and store data as part of Big Data system on centralized distant/global side, with newly-developed challenges.

Development of this concept has evolved in past several years from early proposals like ARM-IBM ARM mbed platform (2014) to new dimension where AI-enabled intelligent devices serve in "enhanced intelligence to IoT applications, meeting the demands of analytics, learning, VR and AR applications in robotics, machine vision and IoT gateways" [5].

Currently, the main promoter of development in this direction is the need for intelligent systems aimed for autonomous vehicle steering. This application of the concept requires robust MCU/SoC systems with high performances, capable of real-time independent work, with a high degree of reliability. If we take the fact that a contemporary car already has about a hundred microcontrollers built-in for current needs, and that the vehicles will need to communicate not only within their own system, but also with other moving cars over the shared link in the near future, it becomes clear why it is necessary to process as much information as possible locally. It will be expected from high-performance MCUs to independently process all their peripheral sensor subsystems and synthetise only important data, which shall be shared with other vehicles and BigData systems. In this way, the global network will be relieved of excessive traffic and amount of data.

The need to enable the vehicle steering systems to 'see' their surrounding through a combination of high-

resolution optical sensors and LIDAR systems can be compared to the living world, where eyes, as the sense of vision, have lead to development of animal species in the early period of evolution. Just as the eyes have required a more complex nervous system, the new optical sensors in sensor systems require the ability to process a volume of information larger by an order of magnitude. Consequently, the existing architecture of sensor systems becomes increasingly complex on a daily basis, giving a new concept to the IoT, currently consisting of sensor systems passively collecting the information and forwarding it via Internet to cloud-centralised Big Data hubs for processing and analytics.

Just like with the animal kingdom, where the emergence of eyes has overrun the ganglion nervous system with multiple nodes in order to develop more complex centralised systems including bigger brains, it is certain that the further development of intelligent sensor systems will be directed towards higher capacity of autonomous processing by so-called peripheral intelligence. At the same time, this is the only way to add adequate performances to future autonomous vehicles. The existing IoT-Big Data system includes the data transfer using asynchronous packet network such as Internet, and their later central data processing in a remote Big Data centre. It is still impossible to enable an adequate response speed in cases where a real-time control is necessary next to monitoring a large volume of sensor information. This practice will be even less possible in the future.

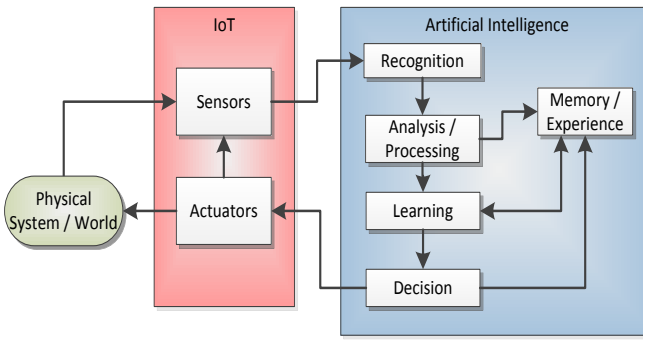

Figure 2. An expanded concept of autonomous AI enabled IoT CPS system (based on the new ARM platform for AI -2016.)

Figure 1 shows the elaborated concept of introducing artificial intelligence to embedded microcotrolling systems, as well as their connection over sensors and initiators with the real physical system. Therefore the need for greater autonomy of processing capacities for local sensor systems and initiators is inevitable. It will be expected that the high-performance microcontrollers intended for work with critical real-time conditions will individually process all their peripheral sensory subsystems by synthesising only a small set of data which will be shared with other vehicles and Big Data Systems either asynchronously or in real time.

This approach will relieve the global network resources from excessive traffic, and Big Data systems from needless preprocessing and compression of excessive information quantities. By relieving the systems of the need to collect more remote system data, which then need to be preprocessed, the microprocessors will preserve the systems' necessary capacities for faster and more demanding analytics of retrieved data, more valuable in meaning.

III. NEW GENERATION OF MCUS

IoT connected device boom on the market in 2014, was followed by the introduction of low-cost ESP8266 WiFi SoC modules based on 32-bit RISC Tensilica Xtensa LX106 MCUs with support for Floating point unit - FPU and Digital signal processing - DSP with base clock of 80-160MHz, and 64KB+96KB RAM. This novelty was full on chip integrated 2.4GHz radio transceiver needed for IEEE802.11n – WiFi. With the price of just a few USD, modules like the one shown in Figure 2 were making their way to all available IoT device applications.

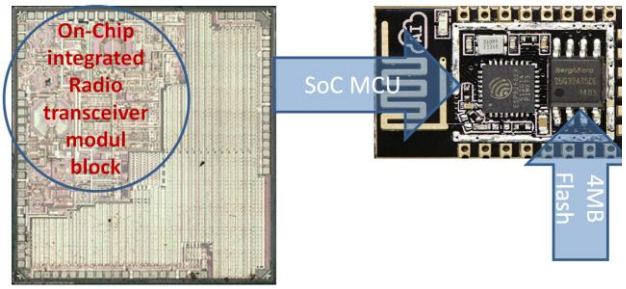

Figure 1. ESP8266 MCU chip SoC arhitecture and typical WiFi enabled module setup with additional 4MB flash storage.

This device marks a distinctive line that separates the new IoT-ready MCUs from other 32bit MCUs, by requiring them to have more than 100MIPS, integrated FPU + DSP, and, most importantly, built-in Ethernet or WiFi support. The rest of ESPs specifications are shown in Table I. Figure 3 presents its typical IoT application.

In the very same year (2014), the market has seen an emergence of the new 32bit generations of IoT-ready microcontrollers with a range of architectures and manufacturers, beginning with ARM Cortex M4 series, based on ARMv7M architecure, introducing an optional FPU-SP unit, ethernet support, MicroCHIP PIC32MZ high-performance MCU series with similar options based on MIPS32 RISC architecture. This first series had a sufficient level of performances for processing the medium-speed sensors and loading local Web servers, enabling performances in the range between 100-200MIPS. The emergence of the second generation of microcontrollers in the end of 2015 and during 2016, being led by ARM Cortex M7, dual core ESP32 and faster PIC32MZ microcontrollers (MIPS32-Warrior) has strengthened these solutions' positions. One shared feature of all these is having built-in FPU+DSP units, LAN support, WiFi (+BT) network connection and significantly more RAM and FLASH ROM neccessary for more complex programmes and multimedia contents. In average, their performances are greater than 200MIPS (32bit) on Dhrystone 2.1 benchmark tests and 2MFLOPS (IEEE 754 64bit-floating point) on Linpack tests.

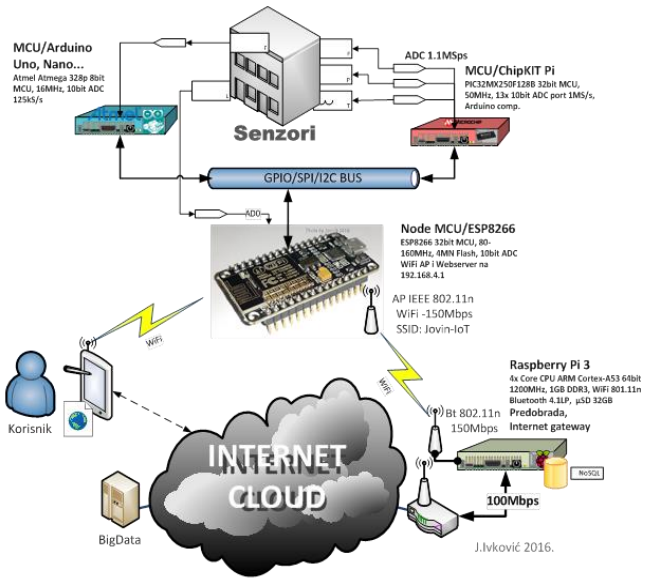

Figure 3. Single MCU (ESP8266)-based solution for control several first gen. 32bit MCUs and provide WiFi connection to Internet [3]

Perhaps the most interesting in this group, ARM Cortex M7 is currently the fastest classical microcontroller solution (nearly reaching lower ARM Cortex R4 Realtime-critical SoC CPU solutions and earlier Cortex – A series). When compared to PC processors, their performances are similar to Intel Pentium 2/3, whereas the energy consumtpion amounts to just a fraction of Pentium's. It is characterised by possessing FPU and DSP.

Apart from the mentioned, the concept of ESP32 SoC MCU is interesting, since it represents the second generation of ESP solutions with WiFi + BT support, but also includes dual core 32-bit RISC Tensilica Xtensa LX106 MCU built-in FPU+DSP unit, increased clock speed (240MHz) and significantly more RAM (520KB). This solution has one core exclusively responsible for network WiFi and Bluetooth connection in real time, whereas the other completely serves for the application support. The manufacturer cites 600 MIPS available on both cores, [6] yet currently, the software development support has limited the performance testing to only the application-assigned core. Therefore, the values represented in Table I are still different.

Moreover, ESP32 MCU can be interconnected into an intelligent mash-network, enabling easier propagation of data as well as solution development, where, if provided adequate application support, the connected sensory network/IoT system can become a self-adapting system (neural network), controlled by its own artificial intelligence (AI). Its processing performances, from the classical viewpoint, would be limited only by the number of nodes and the speed of mutual network communication. These solutions, similar to ARM SoC CPU ones, would be highly energy-efficient, exceeding 400-700 MIPS/W [1].

A. 32-bit microcontrollers – characteristics

ARM Cortex M series with 32bit microcontrollers, based on ARMv7 architecture, were first introduced in 2003. They have enabled use of 16bit 'Thumb' instructions, 4GB linear memory address space, optional 'Memory Protection Unit', priority of interapt and nested interapt-vectors (NVIC). The first presented Cortex – M0/M0+ were memory-saving and optimized to perform dominant 16-bit instructions, with the aim to process and control I/O tasks. Due to their small consumption and higher clock rate, they offered performances greater than 8/16bit MCU companies Atmel and Microchip's solutions. More complex requirements with processing capacities with embedded systems have led to introduction of Cortex –M3 MCU, which have introduced an advanced procession of data, multiple instruction accumulators, higher clock rate, more RAM and instructions enabling manipulation with bit-fields.

In the period between 2004 and 2007, based on ARMv7 and MIPS 4K RISC architecture, a generation of 32bit MCU solution has been introduced, initiating embedded/IoT revolution. Solutions based on ARM Cortex –M0/M0+/M3/M4 and PIC32MX/MZ including solutions from earlier 8/16bit MCU families (Atmel ATMega, PIC10-24 and similar), have initiated embedded computing and the Internet of Things.

In 2013, IoT concept has been introduced as 'Embedded devices with CPU and Memory that can be network connected'. However, a unique definition of it still did not exist.

Apart from ARM Cortex-M series, it is important to point up the Cortex-R series, primarily aimed for high-performance real-time mission-critical embedded computing. Due to insisting on real-time work reliability, this series has primarily found its use in LTE5 telecommunication systems and transport systems (it is now at its version R8).

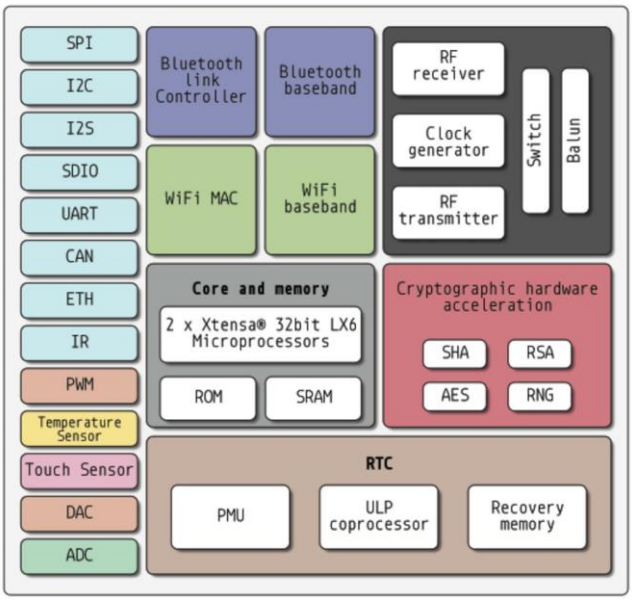

Figure 4. ESP32 dual core "IoT-ready" MCU arhitecture. [6]

TABLE I.
SUMARY TABLE OF MICROCONTROLLER GENERATIONS

NET	Computing power		GPIO ports speed C /Asm.		HW arithmetic acc./ FPU / DSP	Clock / RAM (cash)	CPU/MCU architecture	CPU/MCU	Platforms	Groups by generation
LAN	Floting point instr./sec. **MFLOPS**** Linpack DP	32bit Intiger instr.per sec. **DMIPS*** Dhraystone (C) v2.1 VAX MIPS	Ring-oscilator (2-GIPO's inv. loop)	Bit-banging (I/O toggle rate)						
LAN 100 Mbps	**42**	**875**	8.7MHz	22 MHz	FPU	**700MHz / 512MB**	**ARM1176 (v6) 32bit**	**Broadcom BCM2835**	**Raspberry Pi 1 B+**	**SoC CPU 32-64bit 2012...**
LAN 100 Mbps	**170.92**	**2019.41**	16.92MHz	42.16 MHz	FPU+Neon SIMD	**900MHz / 1GB**	**ARM Cortex-A7 (v7-A) 32bit 4x core**	**Broadcom BCM2836**	**Pi 2**	
LAN 100 Mbps	**180.14**	**3039.87**	19.3MHz	55.8 - 57 MHz	FPU+Neon SIMD+DSP	**1.2GHz / 1GB**	**ARM Cortex-A53 (v8-A) 64bit 4x core**	**Broadcom BCM2837**	**Pi 3**	
LAN 1 Gbps	**885.03**	**8302.97**	9.9MHz	4.. 22.5 MHz	FPU+Neon SIMD+DSP	**2GHz / 2GB**	**ARM Cortex-A15/A7 (v7-A) 32bit 8x core**	**Samsung Exynos5422 Cortex**	**Odroid XU4**	
/	**0.0943**	**10.15 (20)**	1.5894 MHz	3.974MHz	2-cycle Multiplier	**16MHz / 2KB**	**AVR 8bit RISC**	**ATmega328P**	**Arduino Uno / Nano V3**	**MCU 8bit 1996/7**
/	**0.618**	**51.17 (125)**	1.998 MHz	**4.422MHz**	HW int32 MUL/DIV 1-cycle	**84MHz / 96KB**	**ARM Cortex-M3 (ARMv7-M) 32bit**	**Atmel SAM3X8E**	**Arduino DUE**	**MCU 32bit 2004-2007...**
/	**0.369**	**76.28 (83)**	**2.5 MHz**	4.000MHz	HW MUL/DIV 32x32bit-1-cycle	**40-50MHz / 32KB**	**MIPS32 M4K 32bit**	**PIC32MX250 F128xx**	**ChipKit PIC32MX**	
/	**0.552**	**91.75**	723.4kHz	1.5911 MHz	HW int32 MUL/DIV 1-cycle	**72MHz / 20KB**	**ARM Cortex-M3 (ARMv7-M) 32bit**	**STM32F103RB**	**Embedded Pi**	
/	**1.207**	**113.40**	920kHz @160	2.294MHz@160	FPU +DSP	**80-160MHz / 64KB +96KB**	**Tensilica Xtensa LX106 32-bit MCU**	**ESP8266EX**	**Node MCU 1.0**	**MCU IoT Ready 32bit – FPU/DSP 2014...**
/	**2.805 (1x core)**	**176 (224.8)-411.11 on 1-2xcore**	**1.846MHz MHz@240**	***4.543 MHz@240***	FPU +DSP	**160-240MHz / 520KB**	**Tensilica Xtensa LX106 32-bit MCU 2x core**	**ESP32s**	**Node MCU32**	
LAN 100 Mbps	***11.67***	**319 (330 by spec.)**	**8.97MHz**	19.98 MHz	SP/DP FPU +DSP	**200MHz / 512KB (L1 16KB+4KB)**	**MIPS32 Warrior M-class M5150 32bit**	**PIC32MZ2048 EFG100**	**chipKITWiFiRE IoT**	
LAN 100 Mbps	**3.588**	**763.05**	**4.907MHz**	**14.492 MHz**	SP FPU+DSP+Chrom ART Accelerator	**216MHz / 320KB (L1 4KB +4KB)**	**ARM Cortex-M7 (ARMv7E-M) 32bit**	**STM32F746ZGT6**	**NUCLEO - F746ZG**	
LAN 100 Mbps	**22.116**	**744.64**	**4.907MHz**	**14.492 MHz**	DP FPU +DSP +Chrom ART Accelerator	**216 MHz / 512KB (L1 16KB+ 16KB)**	**ARM Cortex-M7 (ARMv7E-M) 32bit**	**STM32F767ZIT6**	**NUCLEO - F767ZI**	

IV. COMPARATIVE TEST AND INDICATORS

A. Synthetic test results

Figures 4 and 5 show the performances of the tested systems in two most significant synthetic tests, Dhraystone 2.1 for integer int32 operations and Linpack test package for floating-point operations (double precision). [7] [8]

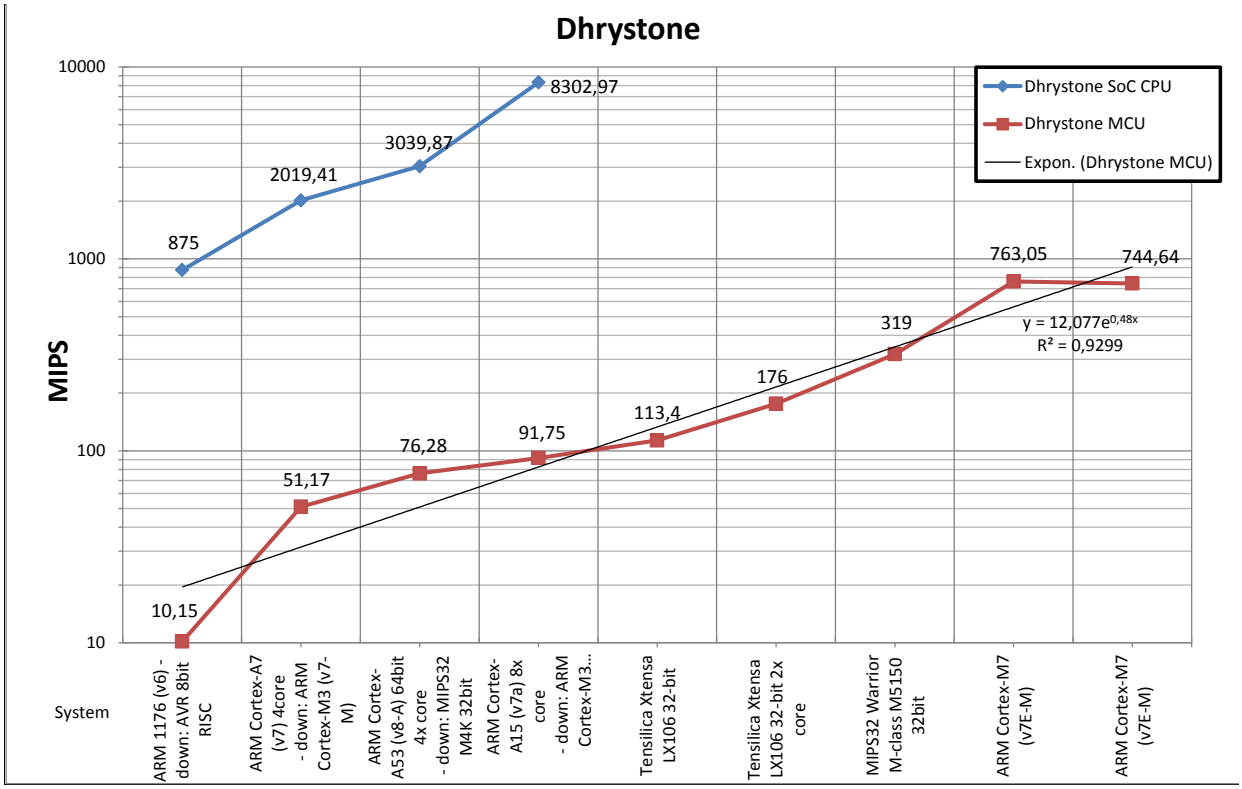

Figure 5. Dhrystone 2.1 benchmark.

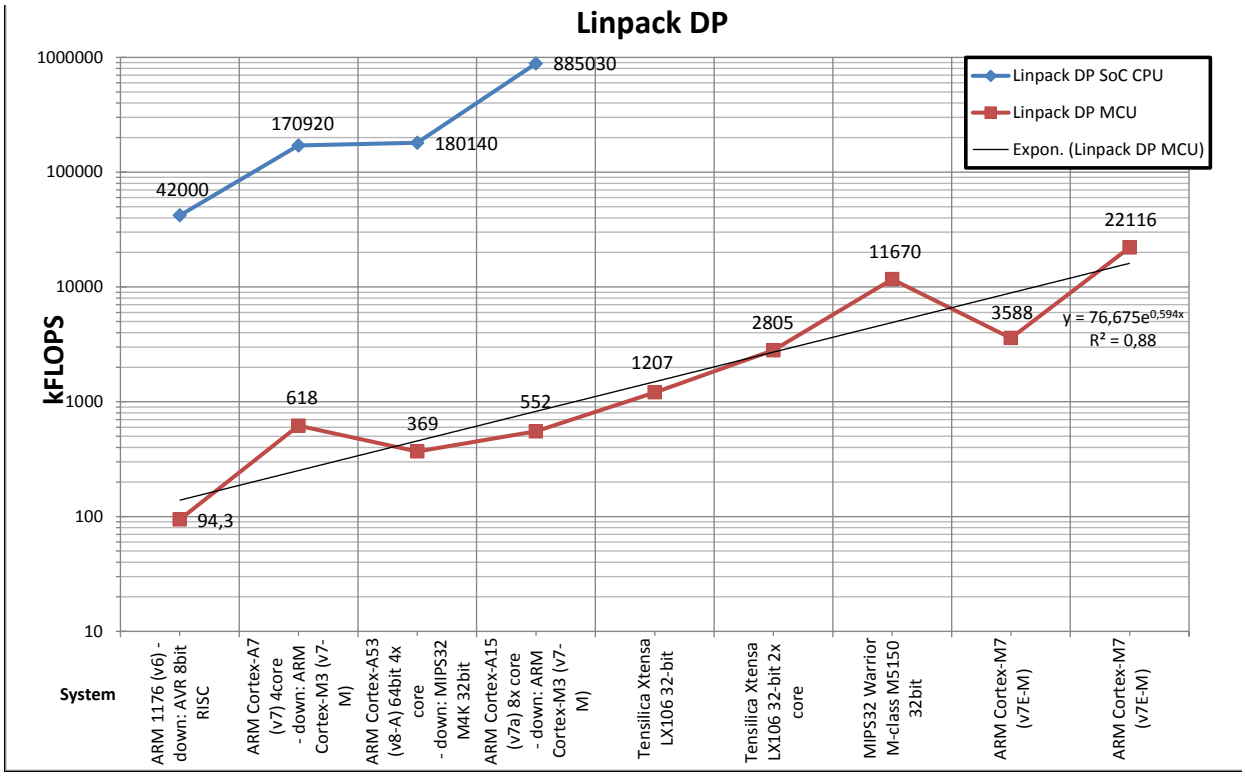

Figure 6. Linpack DP (double precision) FPU benchmark

Figure 6 shows how introducing DP FPU unit in ARM Cortex M7 microcontrollers has improved performances in comparison to SP FPU performance by nearly 5.9 times.

On the other hand, presented growth tendencies are only a roughly set framework, since data on time of certain microcontroller series' occurrence are not inserted. It must be borne in mind that this is a rough comparison of models dominant on market, which necessarily differ in clock speed, RAM memory capacity and other features.

Onwards, the following shall be presented successively: AVR 8bit from 1996. Atmel SAM ARM-Cx.-M3 from 2005, PIC32MX from 2007, ST ARM Cx.-M3 from 2007. These are followed by a new generation with ESP8266-2014, ESP32-2016/2017, PIC32MZ-2012-2014, and ST ARM Cx.-M7 from 2015 and 2016. The time difference between emergence of AVR 8-bit and ARM Cortex M7 32bit microcontrolers is 20 years. It can be noticed that the growth of int32 performances grows by one order of magnitude, i.e. 7 to 10 times per decade, in average. The growth of performances between certain generations within the same architecture like ARMv7 – ARM Cortex M3 and M7 is between 7 and 10 times, whereas somewhat similar example shows two generations of PIC32MX and PIC32MZ MIPS-based MCUs and slightly smaller difference – which in the tested case proved to be about 6 times. However, the fact that the time span between PIC32MX and MZ series' first market appearance was less than five years [9] should be taken into consideration.

This growth ratio is even more striking on the example of difference between ARM Cortex M4 MCU performances (which we were unable to test) with 1.25MIPS/MHz and M7 which has achieved over 3.4MIPS/MHz in our test, where the time span between market appearance of first models M4 and M7 was shorter than 4 years. There is a proof that the performance development and improvement does not stop in the example of an accelerated 'H7' subversion of M7 series, which will bring upon an increase, more precisely doubling of RAM and L1 cache, which will almost certainly double the performances compared to the current state. In this way MCU will achieve performances of Cortex R5 and A7 CPU series, at least in the area of int32 operations with over 1.5GIPS and about 50MFLOPS.

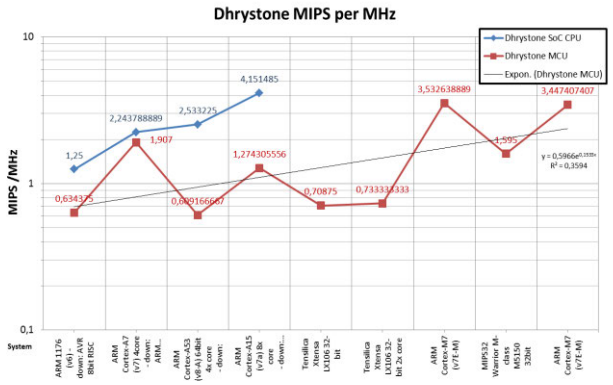

Figure 8. Dhrystone MIPS per MHz rate of MCU architectures.

Figure 7 shows the ratio between Dhrystone MIPS per MHz clock rate and tested MCUs.

B. *Testing of GPIO port performances*

One of the microcontrollers' key features is that they are able to directly monitor and control physical processes thanks to their architecture which includes built-in analog-digital ADC converters, comparators and counters, direct input ports on one and digital and digital-analog DAC outputs on the other hand. Almost all MCU architectures (ESP32 MCU architectures in Figure 4) have this set of structural features in their smaller or bigger scale, but their level of performances is drastically different and easily becomes a limiting factor for certain demanding applications. This analysis gives priority to speed at which the software applications can access and process general input/output ports – GPIO.

When it comes to Bit-banging (I/O toggle rate) test, the speed at which a MCU architecture, using a software code, can control the internal bus, as well as set and erase digital states from a certain port, is crucially significant. This information is different from hardware-solved SPI/I2C protocols and maximum work speed of a port frequently cited by manufacturers. It represents the really achievable performances of software-hardware architecture. A simplified scheme of the test procedure can be seen in Figure 8, implying the use of electric device for measuring the resulting clock rate and the look of the

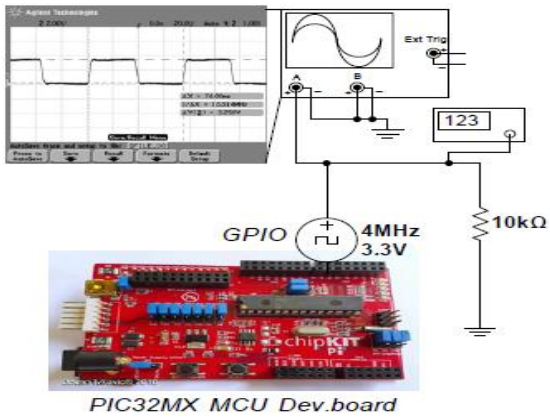

Figure 7. Bit-banging MCU GPIO port test setup.

curve.

The new generation of microcontrollers introduces significant improvements in the area of GPIO port manipulation speed, particularly bearing in mind that the previous generations did not bring upon a significant performance increase even in comparison with a decade older 8-bit MCUs. The growth rate is evident and it can be expected that the PIC32MZ and ARM Cortex-M7 with current 13-20MHz will achieve the level of ARM R and specialised A series in the future iterations.

Another performed test is the Ring-oscillator (2-GPIO inv. loop) which includes a closed loop in which a logical inversion of the output digital port state is done by software. For example, Port_1 related to what is read from it by Port_2 and so forth, successively and cyclically. In this way, a software-defined P1->P2-Inv->P1->P2… circuit is created, which obtained frequency is an excellent indicator of the speed at which MCU can set, read from input port and then return the logically inverted obtained value to the output port.

In both of these tests, a multiple increase of performances in new generation of MCUs is clearly noticeable, compared to the previous generations.

C. Identified trends

After everything mentioned above, an increase in the speed of new MCU generation is visible in the area of analog-digital conversion of input electric units, as well as almost mandatory existance of adequate high-speed DAC output modules. This largely freens the MCU solutions from the need for specialised ADC and DAC circuits, enabling a direct reading of sensory input electric units in this way, as well as controlling analog actuators next to digital and pulse.

A faster bus, peripheral ports and hardware built-in serial SPI/I2C and similar protocols are an inevitable consequence of increase in internal processing performances and sensory environment in which these MCU systems need to work.

V. Conclusion

The results of the measurements carried out show that the latest generation of high-performance MCU platforms themselves or as a part of IoT sensor devices have computing capacity to pre-process complex data locally and to unload the network resources necessary for transmitting large quantities of information for Big Data processing.

If we consider that nature of MCU's work can be considered as pure, somehow parallelized and independent of serialized network exchanges it then has a great potential for Massive parallelized heterogeneous computing in Cyber Physical Systems.

The latest generation of MCUs shows promising increase in performance over previous generations, rate and nature of this increase is exponential and beyond rate proposed by Gordon Moore [10] , but at the same time new MCUs preserve both low consumption and great energy efficiency with inherited robustness and reliability by using older and more mature production technologies. That trend can only accelerate in near future driven by increasing demand for more both powerful and low-cost MCU/SoC devices.

New generation of MCUs can take and perform the role of small on-premises or web servers, that process local sensitive private information (of local importance) while relieving network and Big Data analytics of need to transfer and store great amounts of raw data. In that way, security and confidentiality of locally collected sensitive raw information will be increased, at the same time transferring only a synthetized small subset of data that is of direct interest for particular Big Data analytics.

Acknowledgment

This work was fully financed and supported by authors themselves without any institutional help or support. The author would like to thank the commercial companies: Insel elektronika d.o.o. Novi Sad, Serbia for supporting development with rare electronic components, and Premier Farnell UK Limited, UK for their help in procurement of new MCU/SoC devices.

References

[1] Intel co., "Solution Blueprint Internet of Things (IoT) Improving Manufacturing Processes, Customer Experiences, and Energy Efficiency with the Internet of Things (IoT)", IOT Solution World Congres 2016, Barcelona, Spain, 2016.

[2] J. Ivković, Dissertation: "The Methods and Procedures for Accelerating Operations and Queries in Large Database Systems and Data Warehouse (Big Data Systems)", Zrenjanin: University of Novi Sad, Serbia, 2016, p. 211.

[3] J. Ivković, A. Veljović, B. Radulović and L. Paunović, "4.0 INDUSTRY – SoC/IoT SYSTEMS AND INTERNET INTELIGENT DEVICES", in ITOP2016, Faculty of technical science Čačak, Čačak, Serbia, 2016.

[4] Irena J. Petrick, ""Getting to Industry 4.0 with the Internet of Things", Vol.1 Winter 2016, Intel co.," Intel co., 2016.

[5] ARM Ltd, "Internet of Things," ARM Ltd. UK, 2017.. [Online]. Available: http://www.arm.com/markets/internet-of-things. [Accessed 01 03 2017].

[6] Espressif Systems Inc., China, "ESP32 Datasheet v1.1," 01 03 2017. [Online]. Available: https://espressif.com/sites/default/files/documentation/esp32_datasheet_en.pdf. [Accessed 01 03 2017].

[7] J. Ivković, DHRYSTONE Benchmark C, Version 2.1 for ARM Cortex-M7 MCU in C++ (ST Nucleo-144 Stm32F746 and Stm32F767), London, UK: ARM Ltd. The ARM mbed IoT Device Platform (developer.mbed.org), 2017.

[8] J. Ivković, Linpack port to ARM MCU from Arduino IDE and roylongbottom.org ARM based source, London, UK: ARM Ltd. The ARM mbed IoT Device Platform (developer.mbed.org), 2016.

[9] Microchip Technology Inc.,USA, "PIC32MZ-family: PIC32MZ2048EFG144 Datasheet," Microchip Technology Inc, 11 07 2016. [Online]. Available: http://www.microchip.com/wwwproducts/en/PIC32MZ2048EFG144. [Accessed 20 02 2017].

[10] G. Moore, "Cramming more components onto integrated circuits," Electronics Magazine, pp. 114-117, 19 April 1965.