European Processor Initiative (EPI)
An approach for a future automotive eHPC semiconductor platform

M.Kovač[1], D.Reinhardt[2], O.Jesorsky[3], M. Traub[2], J.M.Denis[4] and P.Notton[4]

Abstract: In this paper we present a novel approach for a future automotive embedded high-performance computing (eHPC) platform. The platform is based on an European Processor Initiative (EPI), Common Platform (CP). This paper also gives an overview of the automotive industry's challenges and the general architectural aspects of EPI CP.

1 Introduction

The importance of high-performance computing has been on the rise in the last years, and it is expected that this trend will not only continue but rapidly grow. Industry reports show that annual global IP traffic will soon reach several zettabytes, that vast amounts of new devices collect and store data, that scientists are exploring new computing approaches to solving global challenges, that the industry is changing the way products are designed, and that we, as individuals, are constantly expecting more personalized services for the entire spectrum of our lives' aspects.

The need to collect and process the vast amounts of data in an efficiently and timely manner comes at a price. The existing approach to HPC systems design is no longer sustainable for the exascale era in which exascale systems are capable of executing 10^{18} calculations per second). Energy efficiency is of enormous importance for sustainability of future exascale HPC systems.

Europe has recognized this global HPC challenge and has strategically initiated the efforts to support the next generation of computing and data infrastructures. EU efforts are synchronized in the establishment of the EuroHPC Joint Undertaking, a legal and funding entity which will enable pooling of the EU and national resources on HPC to acquire, build, and deploy the most powerful supercomputers of the world in Europe.

The European Processor Initiative (EPI) is one of the cornerstones of this EU HPC-strategic plan. EPI draws 23 partners from ten European countries, with the aim to bring a low power microprocessor to the market and to ensure that the key competence of high-end chip design remains in Europe. Thanks to these new European technologies, European scientists and industry will be able to access exceptional levels of energy-efficient computing performance. This EPI efforts will benefit Europe's scientific leadership, industrial competitiveness, engineering skills and know-how, and the society as a whole.

One of EPI's core activities will take place in the automotive sector, providing architectural solutions for a novel embedded high-performance computing (eHPC) platform and also ensuring the overall economic viability of the initiative.

EPI's automotive stream activities are targeting main trends driving the innovations in automotive industry including the introduction of autonomous driving (Levels 4 and 5) and the Connected Car infrastructure. New autonomous vehicle network architectures require computing platforms to be able to execute complex vehicle perception algorithms that include sensor/imaging processing, data fusion, environment sensing and modelling, low-latency deep machine learning for object classification and behavior prediction with seamless, dependable and secure interaction between mobile high-performance embedded computing, and stationary server-based high-performance computing.

[1] University of Zagreb, Faculty of Electrical Engineering and Computing, Zagreb, Croatia
[2] BMW, München, Germany
[3] Elektrobit Automotive GmbH, Erlangen, Germany
[4] ATOS, Paris, France

This is a post-peer-review, pre-copyedit version of an article published in Electronic Components and Systems for Automotive Applications - Proceedings of the 5th CESA Automotive Electronics Congress, Paris, 2018. The final authenticated version is available online at: http://dx.doi.org/ 10.1007/978-3-030-14156-1_15.

2 Future challenges for the automotive industry

The European automotive sector with its original equipment manufacturers and suppliers represents an important share on the global market. In nearly every European country the automotive industry is one of the most important domains for the national economy. Meanwhile, the automotive domain is characterized by high complexity and ICT-driven innovation with steering and control of most of the automotive functions by performant embedded computing units (Electronic Control Unit - ECU).
Applications in the automotive sector have to fulfil highly challenging environmental requirements, endure a typical temperature range from -40 to 125°C and above, as well as mechanical stress and chemical and moisture resistance over a lifetime of at least 15 years. Most of the automotive functions are real-time-relevant and have to fulfil not only functional safety but also cyber-security requirements.

The development in the automotive domain is currently characterized by three main trends:

- The transition from internal combustion engines to full electric vehicles,
- The introduction of highly complex and safety-relevant ADAS assistant systems including first automated driving (AD) functions which demands fail-operational systems, and
- The 'Connected Car', its integration into large-scale traffic management systems, and combining both high-performance embedded computing platforms in cars and exascale-class computing centers.

As a result, automotive innovation trends require significantly increased computing performance. The current vehicle network architectures with up to 100 dedicated ECUs are no longer capable to answer the requirements towards increased computing performance, fail-operational and functional safety, cyber-security, and real-time behavior (RT). To reduce the number of ECUs, large scale software integration (LSSI) is aimed to consolidate smaller ECUs on single integration platforms (Figure 1).

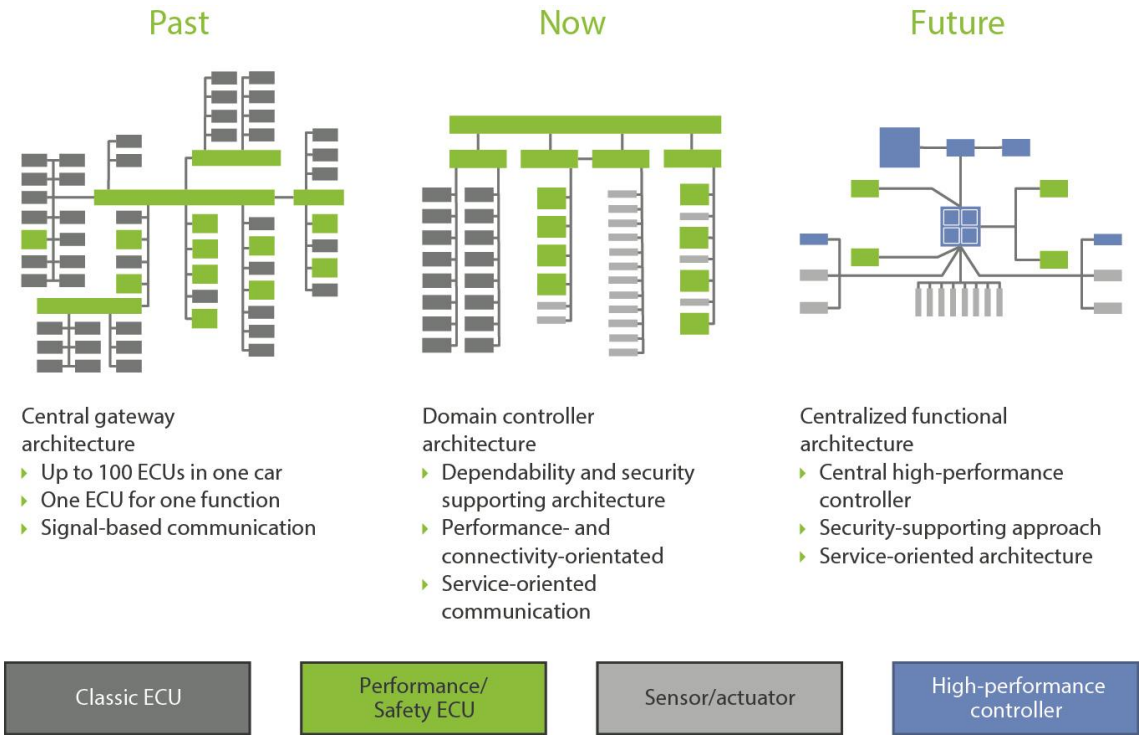

Figure 1: Consolidation of automotive ECUs and functionality towards a centralized function architecture

This is a post-peer-review, pre-copyedit version of an article published in Electronic Components and Systems for Automotive Applications - Proceedings of the 5th CESA Automotive Electronics Congress, Paris, 2018. The final authenticated version is available online at: http://dx.doi.org/ 10.1007/978-3-030-14156-1_15.

As of today, embedded automotive-qualified processors are not powerful enough to handle issues like sensor data fusion from e.g. camera, radar, and lidar. At the moment, the automotive industry is addressing this by using high-performance CPUs and GPUs from other application domains like e.g. NVIDIA and Intel chips from the consumer and performance computer industry in combination with automotive-qualified master ECUs like e.g. the Infineon AURIX.
Specific automotive processors (number crunchers) adapted from the HPC-domain in combination with qualified innovative automotive MCUs could solve the problem. From the system viewpoint, adapting and using consumer processors must under no circumstances have a negative impact on the function, the real-time behavior, and the availability and the reliability of the automotive compute platform.
Future vehicle network architectures that are based on only a few powerful high-performance computing platforms and that are capable to implement vehicle perception tasks in real time in a fail-operational manner are highly required (Figure 1).
Looking to autonomous driving functionality vehicle perception for building a model of the surrounding environment is computationally intensive. It includes sensor processing, computer vision, data fusion, object classification, and behavior prediction. Furthermore, the borders for vehicle network functionality are not fixed and the automotive domains are merging, dissolving, or getting into background for new innovations. Car functionality will be classified in different classes of ECUs [11]. OEM-specific innovations are integrated onto high-performance integration ECUs. These are interconnected by a fast fieldbus like automotive Ethernet.
The expected performance levels for future autonomous vehicles with regards to data produced by the sensors in the vehicles and the bandwidth requirements so the vehicle can properly interact with its environment have been analyzed [*12*] and are shown in Table 1. It can be observed that high performance low power processors and accelerators are a must for Autonomous Driving (AD) Level 4 and Level 5 vehicles.

System Parameters	**Level**		
	3	**4**	**5**
Processing Capacity (TFLOPS)	40	80	120
DRAM (GB)	1	10	20
Non-Volatile Storage (GB)	100	300	500
Data Link Interface (Gbps)	0,1	1	1

Table 1. From AD Level 3 to Level 5 Vehicles: Evolution of processing capacity, storage and data link interface

The anticipated evolution of processing capacity is large (x2) from Level 3 to Level 4 vehicles and from Level 4 to Level 5 (x1.5). In the very classical HPC world, this evolution of performance would be manageable, considering the significant evolution of the electrical power consumption too. But, embedded compute elements have totally different constrains because of the limited embedded electrical power.
Thus, a simplistic approach exclusively based on an evolution of the compute power cannot match the target performance in a reasonable electrical envelope. As opposed to this "*brute force*" approach, the architecture concepts described hereunder will fill the gap between the TFLOPS (10^{18} floating point operations per second) requirements and the electrical power a vehicle can provide.

This is a post-peer-review, pre-copyedit version of an article published in Electronic Components and Systems for Automotive Applications - Proceedings of the 5th CESA Automotive Electronics Congress, Paris, 2018. The final authenticated version is available online at: http://dx.doi.org/ 10.1007/978-3-030-14156-1_15.

3 EPI Common Platform Approach

Exascale computation systems will need to simultaneously meet challenges related to performance, system cost, and energy efficiency. To deliver performance, a vast amount of resources is required, but the wrong choices of components, architecture, or implementation might result in a system which is much too expensive and too power-hungry. To find the right balance, global system level optimization is necessary.
EPI therefore intends to share a strong set of common technology between different application domains. Starting from the selection of cutting-edge processor technology, a low-power design approach will be centralized in the hardware around massive parallelism, specialized architecture, low-voltage operating point and fine grain power management. The software stack will be designed to integrate and take advantage of these features to achieve high-energy efficiency and maximize performance across a wide range of layers from the low level firmware, all the way up to system software and application run-times. For this purpose, EPI will harmonize the heterogeneous computing environment by defining a common approach: the so-called **Common Platform (CP)**. It will include the global architecture (hardware and software) specification, common design methodology, and global approach for power management and security.
The CP is organized around a 2D-mesh Network-on-Chip (NoC) connecting computing tiles based on –a general-purpose CPU core with built-in FPU acceleration (e.g. 64b core with vector engine), RISC-V based high energy efficiency accelerators with different acceleration levels or any other application-specific accelerator.
A common software environment between heterogeneous computing tiles will harmonize the system as well as acting as a common backbone of IP components for IO connection with the external environment such as memories and interconnected or loosely coupled accelerators. With this CP approach, EPI will provide an environment that seamlessly integrates any computing tile. The right balance of computing resources for application matching will be defined through the ratio of accelerator and general-purpose tiles.

4 Concepts for automotive HPC Platforms within EPI

One of the key challenges of EPI is to explore and validate the scalability and deployment of computation concepts from HPC down to embedded application domains like automotive. In the following, we describe our methodology for automotive eHPC hardware and software codesign.

4.1 Requirements for the automotive eHPC

In order to focus on sustainable concepts for embedded computing platforms capable of real-time processing, it is evidently necessary to improve architecture, connectivity, and computing performance of automotive-qualified multi-cores. The outcome of the EPI in this area should be a performant fail-operational automotive real-time multi-core processor.

4.1.1 Processing power

Future emerging automotive applications raise issues like big data handling, deep machine learning, and combined on-board processing with server-based offloading, which will strongly rely on a seamless and dependable interaction between mobile high-performance embedded computing and stationary server-based high-performance computing. From this perspective, a lot of similarities are visible and therefore opportunities for collaboration between automotive and HPC in terms of architecture, hardware, and software arise. Within the scope of the EPI, it is the intention to explore and validate the scalability and deployment of computation concepts from HPC down to embedded application domains like the automotive domain.

This is a post-peer-review, pre-copyedit version of an article published in Electronic Components and Systems for Automotive Applications - Proceedings of the 5th CESA Automotive Electronics Congress, Paris, 2018. The final authenticated version is available online at: http://dx.doi.org/ 10.1007/978-3-030-14156-1_15.

4.1.2 Real-time performance

In order to focus on sustainable concepts for embedded computing platforms capable of real-time processing, it is also evident and necessary to improve the architecture, connectivity, and computing performance of automotive-qualified multi-cores. The outcome of the EPI in this area will be a performant fail-operational automotive real-time multi-core processor.

Apart from the automotive domain, this type of platform might be used as a reference implementation for other high-performance embedded computing domains like industry, machinery, and healthcare.

4.1.3 Environmental settings

The intended system Automotive eHPC Compute Platform must operate in a typical ambient temperature range from -40°C to 125°C and beyond. In addition, defined usage profiles have to be observed. In any case, the automotive-qualified processors (eHPC-MCU) must be able to monitor the correct function of the slave processors. They must safeguard the real-time capability, protect the system for security reasons, and perform a reduced application in parallel in order to take over in case of unavailability or degradation of the HPC GPP.

4.1.4 Energy efficiency

The trend for future vehicles within the next decade is strongly pushing towards electrical cars or hybrid variants. Therefore, energy-efficient chips are necessary to fulfill the previously mentioned environmental requirements for heat dissemination, chip robustness, and reliability. But similar to smartphones, the computational outcome (MIPS per Watt) must be taken into account for future HPC platforms. The goal should be to extend the electric operation and the resulting driving range (battery-driven) for as long as possible.

4.1.5 Safety

Autonomous driving is mentioned as one of the most customer-intensive game changer of the next decade. This freedom brings new requirements for safety and security. The driver will not be the main controlling instance in severe traffic situations. As a consequence, the eHPC computing platform has to achieve fail-operational requirements in order to be suitable for VDA (German Association of the Automotive Industry) Autonomous Driving Level 4 and 5.

4.1.6 Security

Because cars have typically a long product lifecycle time, the possibility for updates and even upgrades are necessary in our next car generations. Nowadays, security mechanisms can be obsolete in five or more years. Computing power will rise more and more and today's encryption mechanisms could be broken in the future. Looking to quantum computing, even new technologies must be taken into account, which brings new possibilities and dangers in parallel.

4.2 Hardware platform for the automotive eHPC

Without innovative solutions, the digital progress in the automotive sector will end in a deadlock because of insufficient computing power for new and increasing fields of application like 360-degree environment recognition and other real-time systems.

In order to continue on the road of success and to overcome this dilemma, the basic approach is the use of suitable high-performance processors from the HPC sector. The specific challenge is the integration of those high-performance consumer processors that have to be compliant with the relevant application and environmental requirements in the real-time domain of automotive.

This is a post-peer-review, pre-copyedit version of an article published in Electronic Components and Systems for Automotive Applications - Proceedings of the 5th CESA Automotive Electronics Congress, Paris, 2018. The final authenticated version is available online at: http://dx.doi.org/ 10.1007/978-3-030-14156-1_15.

A powerful data fusion platform, such as the automotive embedded HPC platform, could then be obtained by combining such automotive-certified real-time multi-core processor with adapted HPC-General Purpose processors and accelerators. Apart from the automotive domain, such an eHPC platform might be used as reference implementation for other safety-oriented, eHPC-computing domains like industry, machinery, and healthcare.
Necessary improvements and extensions of the automotive multi-core architecture and fast interfaces to the accelerator and the high-performance computing general purpose processor (HPC-GPP) have to be considered as well as suitable software concepts and complex automotive application scenarios using on-board and off-board computation (e.g. traffic management, AD systems with machine learning). This type of interaction will require reliable and secure communication channels, proper identity management and assurance, while providing adequate data and identity privacy. Therefore, it is necessary to integrate hardware/software co-design but also safety/security co-design to successfully implement an automotive eHPC platform.
The favored fundamental approach is a master-slave constellation with one or several automotive-qualified processors operating as masters and one or several HPC-GPPs and accelerators operating as slaves.
With respect to the HPC domain, one of the main incentive reasons to adapt the HPC-GPP - not only for the original purpose with local servers and super computers but also for embedded domains like automotive - is the expected improved exploitation of the highly expensive development of such a complex GPP.
From the automotive viewpoint, the expected additional computing power should not exceed given cost and energy budgets.
A proposed subsequent integration of the HPC-GPP and HPC accelerators into the automotive eHPC is supposed to be technologically, functionally, and economically successful. In order to ensure this, both from the perspective of the HPC world and the automotive point of view, a couple of fundamental questions must be determined and clarified right from the beginning.

4.3 Software platform for the automotive eHPC

As described in the previous chapter, a master-slave constellation with one or several automotive-qualified processors operating as masters and one or several HPC-GPPs and accelerators operating as slaves is targeted.
In order to make best possible use of this powerful architecture, the corresponding software requires a more flexible and suitable architecture, which can represent these dynamics.
For EPI, a Classic AUTOSAR for the master controller will be supplemented with an Adaptive AUTOSAR platform running on the HPC-GPPs. This enables a dynamic software configuration and, with service-based communication and heterogeneous computation, it provides mechanisms that ensure the necessary performance. The hardware resources of the performance cores are separated by a hypervisor. The hypervisor virtualizes the hardware and, in so doing, provides the partitions as virtual machines. In this way, the various Adaptive AUTOSAR and, if required, also Classic AUTOSAR partitions can be created.
The combination of a Classic AUTOSAR system with Adaptive AUTOSAR is a state-of-the-art approach to tackle the combination of performance and security requirements of the vehicle network architecture for technologies like autonomous driving. Figure 2. provides a sample configuration of a future automotive multi-core performance controller.

This is a post-peer-review, pre-copyedit version of an article published in Electronic Components and Systems for Automotive Applications - Proceedings of the 5th CESA Automotive Electronics Congress, Paris, 2018. The final authenticated version is available online at: http://dx.doi.org/ 10.1007/978-3-030-14156-1_15.

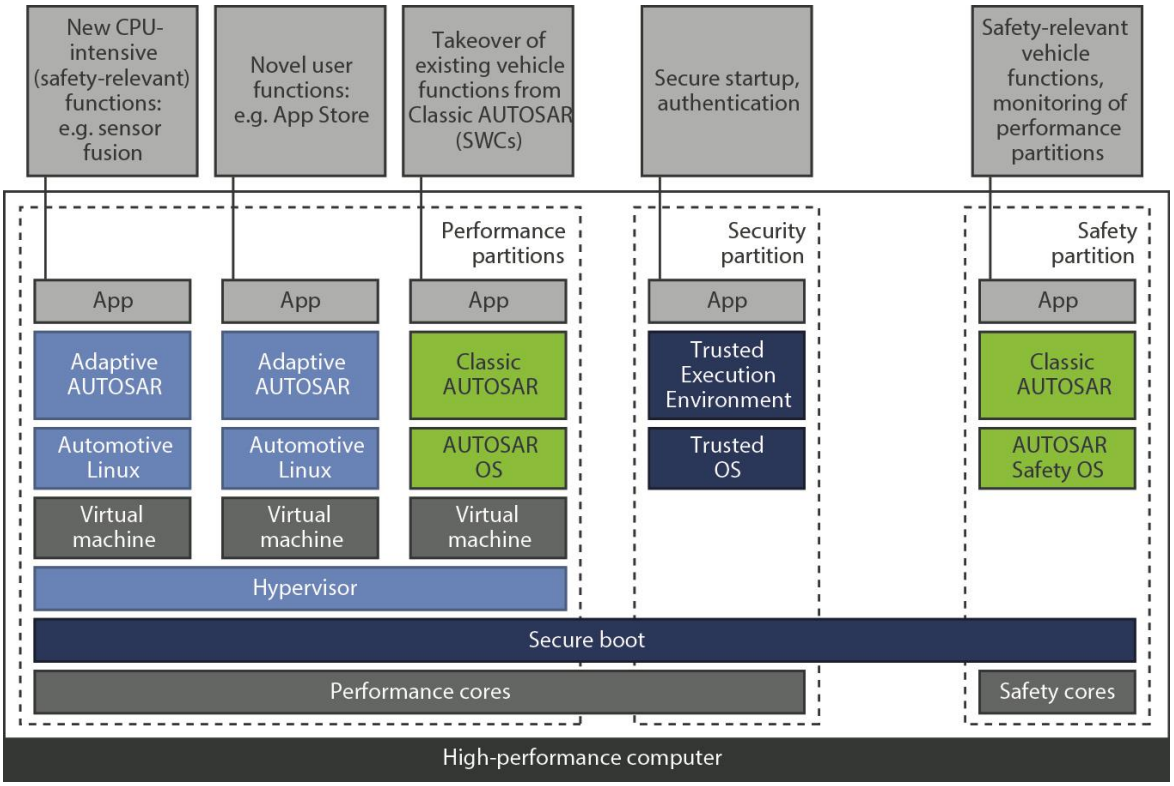

Figure 2: Architecture of a central control unit with performance and safety cores

The performance of the EPI eHPC platform will be measured with real-world automated driving use cases. Functional modules with high computational demands, such as sensor data fusion and AI components, will be executed in EB robinos [9], a software framework for automated driving which is running on top of Adaptive AUTOSAR.
Code optimization of the functional modules will guarantee the best possible use of the specific EPI platform capabilities, e.g. accelerators. Measurement results will be benchmarked with already existing automotive high-performance controllers.

5 Summary

The European automotive sector with its original equipment manufacturers and suppliers represents an important share on the global market. The automotive domain is characterized by high complexity and ICT-driven innovation with steering and control of most of the automotive functions by performant embedded computing units. In order to focus on sustainable concepts for embedded computing platforms capable of real-time processing, it is evidently necessary to improve architecture, connectivity, and computing performance of automotive qualified multi-cores.
The importance of general high-performance computing has been on the rise in the last years, and it is expected that this trend will continue to grow. Europe has recognized this global HPC challenge and has strategically initiated the efforts to support the next generation of computing and data infrastructures. The EPI is one of the cornerstones of this EU HPC-strategic plan.
One of EPI's core activities will be in the automotive sector, providing architectural solutions for a novel eHPC platform and also ensuring the overall economic viability of the initiative. EPI will harmonize the heterogeneous computing environment by defining a common approach: the so-called Common Platform (CP). It will include the global architecture (hardware and software) specification, common design methodology, global approach for

This is a post-peer-review, pre-copyedit version of an article published in Electronic Components and Systems for Automotive Applications - Proceedings of the 5th CESA Automotive Electronics Congress, Paris, 2018. The final authenticated version is available online at: http://dx.doi.org/ 10.1007/978-3-030-14156-1_15.

power management and security. A powerful data fusion platform, the automotive eHPC platform, will be obtained by combining automotive-certified real-time multi-core processor with novel EPI HPC processor that includes accelerators customized to automotive domain. We consider this to be winning approach for a Future Automotive eHPC Semiconductor Platform.

Acknowledgment

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 826647.

References

1. Dr. Gentner, H.; Manz, M.: *"Environment Model for Autonomous Driving – The upcoming Challenges".* ConCarExpo Conference, BMW AG. 2017. Berlin
2. J. Flich, M.Kovac et al: "Enabling HPC for QoS-sensitive applications: the MANGO approach", Proceedings of the Design, Automation and Test in Europe (DATE) 2016, pp. 702-707.
3. Albrecht Meyer; Karsten Schmidt; Denny Marx; Jens Harnisch; Herbert Christlbauer: "Non-intrusive Tracing at First Instruction"; SAE 2015 World Congress&Exhibition; DOI:10.4271/2015-01-0176, 2015.
4. Jens Harnisch; Rolf Schneider; Andre Kohn; Karsten Schmidt; Sven Schoenberg; Qian Zhou:"Efficient Virtualisation for Functional Integration on Modern Microcontrollers in Safety-Relevant Domains"; Article, SAE Technical Papers, Apr 2014.
5. Albrecht Meyer; Peter Gliwa: "Measurement of application performance on Automotive microprocessors"; VDI Wissensforum "Elektronik im Kraftfahrzeug"; 2011, Baden-Baden.
6. Christian Bradatsch, Theo Ungerer, Rafael Zalman, Andre Lajtkep: "Towards Runtime Testing in Automotive Embedded Systems"; Proceedings of the 6th IEEE International Symposium on Industrial Embedded Systems (SIES), pages 55-58
7. Ovidiu Vermesan, Mariano Sans, Peter Hank, Glenn Farrall: "Advanced Electronic Architecture Design for Next Electric Vehicle Generation",Lecture Notes in Mobility, ISSN 2196-5544, 2015
8. Roman Pallierer, Börge Schmelz: "Adaptive AUTOSAR for high-performance in-car computer", EB TechPaper, 2017, https://www.elektrobit.com/tech-corner
9. Dr.-Ing. Sebastian Ohl: "Staying in lane on highways with EB robinos", EB tech paper, 2017, https://www.elektrobit.com/tech-corner
10. Rudolf Grave: "The vehicle architecture of automated driving", EB tech paper, 2017, https://www.elektrobit.com/tech-corner
11. Traub, M; Maier, A; Barbehön, K; "Future Automotive Architecture and the Impact of IT Trends"; Journal of Intelligent Systems; 0740-7459/17; IEEE; 2017
12. Gartner, "Hype Cycle for Automotive Electronics, 2018 ", 25 July 2018, https://www.gartner.com/doc/3883996/hype-cycle-automotive-electronics-
13. Mario Kovač, "Changing Approaches to HPC Systems", CIO Applications Europe, 2018, https://www.cioapplicationseurope.com/cxoinsights/changing-approaches-to-hpc-systems-nid-520.html