Editorial

Welcome to the 4th EMC² Newsletter which will provide an impression on selected EMC² topics and intermediate project results. In this issue we will focus on:

I. Joint Project Conference and Exhibition with French Systematic Cluster 1
II. EMC2 design environment integration for mixed-critical CPS 2
III. Mechanisms for dynamic runtime environments and services - highlights and cooperation 5
IV. Automotive Applications with Service Oriented Architecture (SOA) 9
V. EMC² use case "Highly automated driving" 11
VI. EMC² use case: "Design and Validation of Next Generation Hybrid Powertrain" 14
VII. COTS Multicores in Avionic Applications 15
VIII. FPGA PARTIAL RECONFIGURATION IN SPACE APPLICATIONS 17
IX. EMC² "Internet of Things" use cases 19
X. EMC² use case: "Video Surveillance for Critical Infrastructure" 21
XI. Developing safety & security co-engineering methods for cyber-physical systems – Windows of Opportunity for Standardization 22
XII. EMC² Workshops and Special Sessions at renowned Conferences in 2016: CPS-Week, SAFECOMP, Euromicro DSD/SEAA 27

I. Joint Project Conference and Exhibition with French Systematic Cluster

The 2016 EMC² project conference was organized together with the Systematic Cluster following the main goals:

- Collect feedback of different users.
- Share experiences and best practices.
- Inform the French ecosystem on the guidelines of the project.

The joint event took place on Sept. 28 at Préfecture des Paris et d'Île-de-France in Paris. The morning session was organized by the Systematic Cluster by four working groups: Digital Trust & Security, Systems Design & Development tools, Automotive & Transport and Smart Cities.

Video recordings of all presentations are available via Youtube:

- Introduction and Agenda (Gérard Poirier, Dassault Aviation)
- EMC2 (presentation of Werner Weber, Infineon)
- Embedded France: The French Embedded Cluster (Cedric Demeure, VP at Thales & President of Embedded France)
- S3P project (presentation of Eric Bantegnie – Leader of S3P project – VP at Ansys)
- Innovation through collaborative projects : From an innovation project to a product, software, digital technologies and business
- Presentation of Proxima et Contrex projects

Impressions from the EMC² conference and exhibition, Sept. 28, 2016, Paris

The afternoon session was organized by EMC² building a bridge to other related projects. Both sessions were accompanied by EMC² demonstrations showing best practices in mixed-criticality, multi-core embedded

systems design and realization. More than 200 persons visited the demonstrations at the conference. In addition, the presentation sessions were made available remotely, where further 158 persons watched either the live streaming or the video on Youtube.

Presentations:

- EMC² - A platform project on embedded micro-controllers in applications of mobility, industry and the internet of things (plenary talk, Werner Weber, Infineon Technologies)
- The S3P project - Smart, Safe and Secure Platform (Eric Bantegnie, Ansys)
- PROXIMA - Improving measurement-based timing analysis through randomisation and probabilistic analysis (Francisco J. Cazorla, Barcelona Supercomputing Center)
- CONTREX - Design of embedded mixed-criticality CONTRol systems under consideration of EXtra-functional properties (Kim Grüttner, OFFIS)

Demonstrations:

- art2kitekt - A modelling and analysis tool for aerospace domain (Sergio Sáez, ITI)
- ZG3D - Parallelization of an industrial inspection system (Juan Carlos Pérez, ITI)
- Multicore application design using embedded Procedure Call Library (eRPC) (Petr Lukas, Marek Novak, NXP CZ)
- A Native segregation of multiple virtual networks over only one physical link (Philippe Ravier, P. Aristote, SILKAN)
- A Benes Based NoC Switching Architecture for Mixed Criticality Embedded Systems (Steve Kerrison, UoBR)
- A Hardware Platform for distributed Radar Processing (Martin Terry, Zuhal Clarke, IFXUK)
- A Safe, Secure and Adaptive Mixed-Criticality System (Youssef Zaki, Detlef Scholle, ALTEN)
- What is this thing called OSLC? (Johnny Öberg, KTH)
- Modelling Support for a Linked Data Approach to Tool Interoperability (Jad El-khoury, KTH)
- Space Platform Applications (Dario Pascucci, TASI); demonstration
- m2cpp – exploring multi-cores by generating c++ from MATLAB (Hans Petter Dahle, Fornebu)
- Object detection on FPGA : Signal, image and video processing (Marin Musil, Petr Musil, BUT)
- Interference Measurement and Mitigation Means on Multi-Core COTS processors (Jimmy Le Rhun, Thales)
- Blind hypervision to protect virtual machines (Olivier Heron, CEA); demonstration
- Many-Core Architecture Integration Experiments in a Representative Avionics Environment (Moha Ait Hmid, CEA)
- Safe Dynamic Resource Allocation in Mixed-Critical Systems (Paul Debrulle, CEA)
- Integrated Specification, Design and Documentation flow for Mixed Criticality Embedded Systems (Pfeiffer, Magillem)
- VITRO - Vision Testing for Robustness (Oliver Zendel, AIT)
- Virtualization-based security and fault tolerance (Zhendong Ma, AIT)
- WEFACT – System qualification and certification (Erwin Schoitsch, AIT)

II. EMC2 DESIGN ENVIRONMENT INTEGRATION FOR MIXED-CRITICAL CPS

Frank Oppenheimer OFFIS

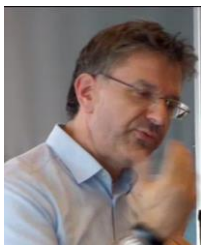

Mladen Berekovic TU Braunschweig

Adam Kostrzewa TU Braunschweig

Haris Isakovic, TU Vienna

The concept of cyber-physical system has been introduced to unite and bind all individual disciplines involved in the whole process from early design to product implementation of industrial or consumer applications. The concept connects the digital and the physical world through a series of mutually dependent interdisciplinary steps.

Once completely independent system components are now co-designed, modeled and simulated to increase the overall performance and efficiency. In EMC², we are covering the whole spectrum of interdisciplinary activities in the CPS design process, from specification in WP1 to the mechanical actors interacting with the physical environment in different use cases (WP7-WP12).

The efficient development of mixed-critical CPS requires a seamless interplay between different layers in the design stack. The hardware layer must provide a reliable electronic architecture supported through services in the run-time system that provide a hardware abstraction layer while allowing to segregate functional and extra-functional behavior. Based on these mechanisms, the application layer must support the explicit modelling of the intended behavior including its criticalities and communication interrelations.

In EMC² these layers are being addressed in three corresponding workpackages (WP). WP2 provides modelling and design languages at application layer, which allow the designers to specify the application tasks including their timing behavior, communication relations and different criticalities. Tools for analyzing, mapping and optimizing the scheduling of mixed critical applications support the designers. With the support from an efficient run-time environment, applications can be executed on modern hardware platforms. The temporal and spatial segregation of different tasks needs sophisticated run-time mechanisms.

WP3 develops a large set of such technologies supporting various advanced hardware platforms. The implementation of mixed-critical applications on multi-core hardware platforms is a specific challenge. Only when the hardware platform provides support for virtualization and predictable communication and timing behavior, the run-time system and thus the application can be implemented according to its requirements.

In order to enhance mechanisms and architectures for the run-time environments (RTE) in WP3, we firstly conducted evaluation of the existing platforms with respect to the support for mixed-critical systems, security techniques as well as safety and real-time properties. Later technologies targeting WP3 goals were proposed including mechanisms for virtualization, hypervisors and monitoring. The main objective was to adapt existing systems to the increasing application system dynamics without losing effectiveness or efficiency. Partners considered on-chip architectures as well as off-chip solutions. Consequently, selected platforms from WP4 were analyzed in WP3 with respect to the dynamic support for Quality-of-Service guarantees targeting different aspects of the design as interconnect communication or resource allocation. For instance, we analyzed the IDA-NoC belonging to the SoCRocket framework (WP4) and proposed the control layer allowing decoupling mechanisms for admission control from the data flow control in the system. Similarly, we proposed mechanisms for the dynamic re-configuration for on-chip networks applicable, which is the part of the Arria V SoC platform (WP4). The proposed extensions were evaluated with respect to the safety requirements, overheads and possible performance benefits. One of the evaluation goals was to build a basis for future extensions of mapping and analysis tools considered in WP2, e.g. art2kitect, to allow easier and flexible system design and implementation which could support feasibility check for a particular mechanism on a selected platform.

WP4 explores and extends existing hardware architectures, concepts and platforms. It covers different levels of abstraction in communication, from an on on-chip and inter-chip communication, to a component to component communication, or a system to system communication. Moreover, WP4 is working on integration of peripheral devices, validation and verification, and dynamic reconfiguration capabilities in hardware.

In the design process for CPS, a hardware platform is often co-designed with the rest of the system. This enables a tighter integration between individual components. A hardware platform is a set of hardware components connected together such that they provide a specific set of functionalities. It can be implemented completely using COTS components (e.g. EMC2-DP, XILINX ZYNC, AURIX), or custom made using programmable logic (e.g., SoC Rocket, Time-triggered MPSoC, MCENoC). The advantages and flaws of one or the other approach depend on the specific set of requirements for each application. These are not always exclusively related to functional requirements of the system, but also with other conditions i.e., certification, cost, or time to market. The design process is highly

influenced by instruments and tools used in the process and these have a huge impact on system integration. A standalone hardware platform without an efficient run-time environment or software modeling approach is a small part in a highly complex puzzle (Figure 1). This is why the unified design and implementation process between different disciplines in both digital and physical part of the system is necessary.

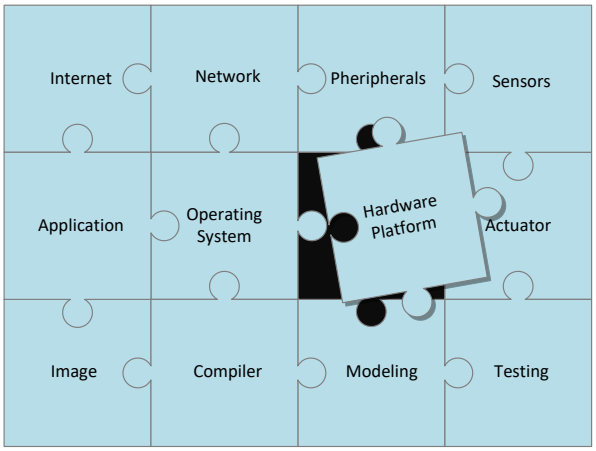

Figure 1: The design environment „puzzle"

Tools for hardware design and configuration need to reach out towards tools and mechanisms on higher levels of abstraction (e.g., on application layer) and provide information necessary for seamless integration. A tool like this can be used to map and optimize applications on specific platforms. Furthermore, integration of custom hardware using programmable logic is more complex and requires significantly more design efforts. We are identifying tools and mechanisms that are able to bridge this gap and provide efficient way to use custom hardware in a domain and application independent fashion.

During the first half of the EMC2 project, we were focusing on developing the necessary modeling languages, mechanisms and technologies. In the second half, the attention shifts more towards the integration and interplay of these technology building blocks. Rather than keeping this effort on a theoretical basis, we decided to aim at the practical integration of at least some major results of WP2, WP3 and WP4. Figure 2 depicts a design environment, which contains individual contributions from all three workpackages. While every design environment must provide solutions to all relevant activities, the particular needed technologies depend on the requirements of the use-case application.

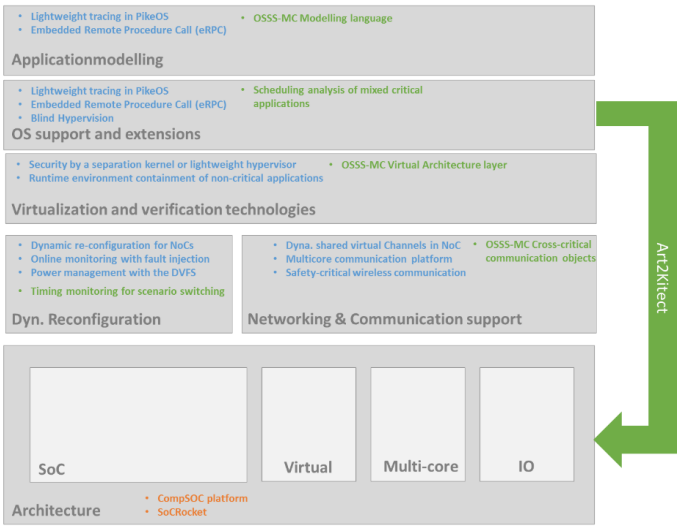

Figure 2: Practical example of a design environment

Without a practical example, this approach would remain a theoretical concept with no true evidence of being applicable in practical product development. In order reduce the significant design effort of such a use-case study, we have chosen the well-understood internal use-case from WP2. The multi-rotor application contains all significant components of a mixed-critical application and is available in full source code. There is already reference implementation using significant parts of WP2 modelling approach and has been implemented using a COTS (XILINX ZYNC SoC) hardware architecture and standard run-time systems.

We are currently applying the major WP2 mapping and analysis framework (art2kitect). This allows modeling the application tasks with its abstract timing requirements and a rather coarse grain representation of potential hardware platform alternatives (see Figure 3).

The aim is here to model different variants of hardware platforms developed in WP4, which are supported by mechanisms coming from WP3. We want to demonstrate that our design environment can address those different hardware platforms (see Figure 4) as developed in WP4.

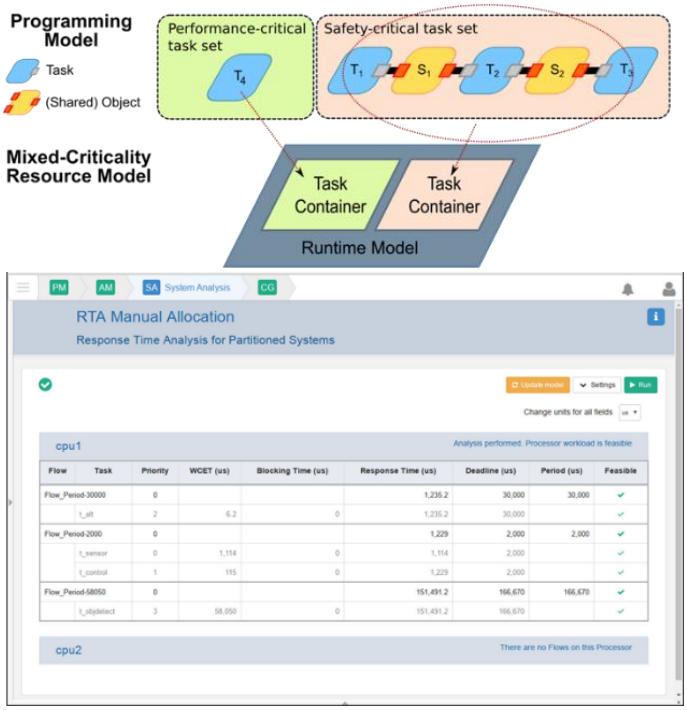

Figure 3: Top: application Modell in OSSS-MC, Bottom: Timing analysis in art2kitect

In the design process, we want to prove that the application model can be developed independent of the hardware platform using appropriate run-time environment mechanisms. The main achievement of this activity will be to show the integration of contributions from three different workpackages into a seamless design environment for a complete application.

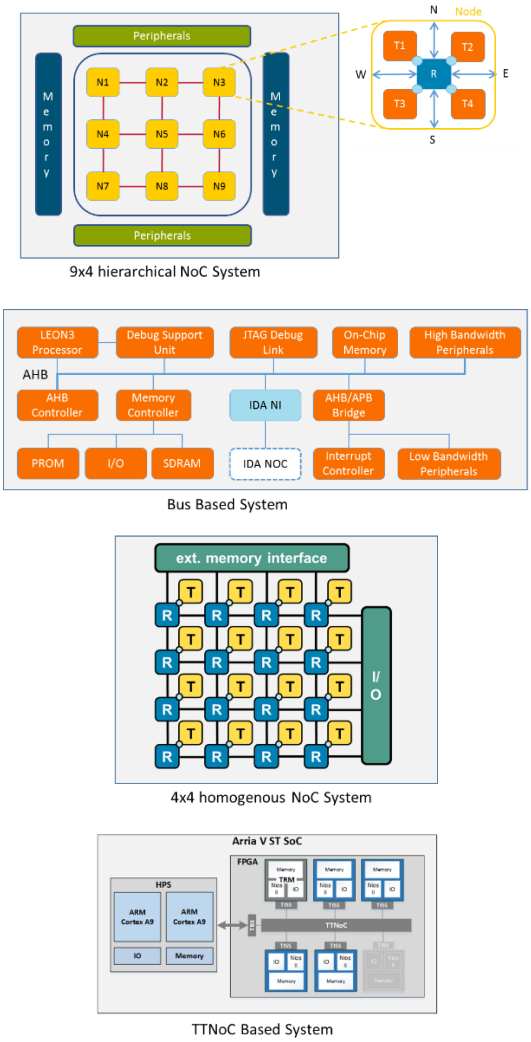

Figure 4: Architecture alternatives

III. MECHANISMS FOR DYNAMIC RUNTIME ENVIRONMENTS AND SERVICES - HIGHLIGHTS AND COOPERATION

Adam Kostrzewa, TU Braunschweig

The broad objective of WP3 is to enhance the European knowledge in mechanisms and architectures for the existing run-time environments (RTE) that are able to support mixed-critical systems, security techniques, safety and real-time properties. Technologies include mechanisms such as virtualization, hypervision and monitoring, which need to be adapted to the increasing application system dynamics with no loss in effectiveness and minimal loss in efficiency. In this article, we present the broad overview of the conducted work using as examples demonstrators presented during the review meetings and conferences. Additionally, we show from the perspective of WP3 the technology transfer and cooperation with WP4 ("Multi-core hardware architectures and concepts") and WP2 ("Executable Application Models and Design Tools for Mixed-Critical, Multi-Core Embedded Systems").

Among the main challenges of the WP, the most important one was to propose solutions applicable to different RTEs, i.e. solutions that are independent from

specific hardware and communication networks and, instead, rely on well-defined interfaces. Although the existing RTEs are domain specific, the underlying principles show many similarities. Therefore, in the EMC² project, partners of WP3 exploit these similarities to simplify portability and safety, as well as cross domain usage and integration of systems, including technology transfers between WP2, WP3 and WP4. In the last four deliverables, which were submitted in the scope of WP3, almost 80 inputs were gathered proposing nearly 40 mechanisms and improvements. New RTE functions are developed and deployed driven by the strategy and methodology based on the evaluation, cooperation and validation.

WP3 is split into six technology clusters grouping the mechanisms in different tasks based on their purpose and scope: communication services (T3.2), virtualization and isolation (T3.3), security mechanisms and services (T3.4), safety platform and real-time mechanisms (T3.5), dynamic application and platform control (T3.6) and OS support functions (T3.7). The coverage of the underlying system architecture by different technology clusters is presented in Figure 5, which depicts a broad overview of the undertaken effort.

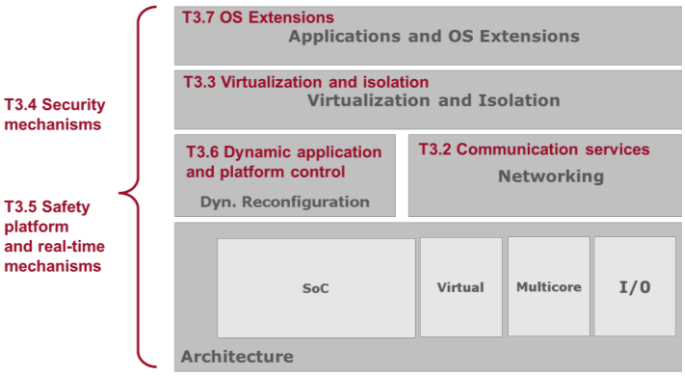

Figure 5: Connection between tasks in WP3 and the SoC architecture originating from WP4

Some of the clusters are targeting specific layers of the system, whereas others propose inter-layer solutions. Moreover, many of the introduced mechanisms offer complex and mature solutions with particular elements (aspects) developed in the independent tasks. The general architecture chart follows the structure proposed in WP4, for easy mutual identification of concrete problems and solutions in both WPs as well as direct and indirect partner cooperation. The direct cooperation considers transfer and application of the mechanisms in architectures from WP4, as done for instance by TTTech and their mechanism for dynamic network re-configuration in. Similarly, Infineon proposed the microcontroller firmware compliant with standard ISO26262 for fast and bug-free development targeting the AURIX chip, applied in WP4 for failure safe operation and simplified development. The indirect cooperation considers the information exchange as well as evaluation of application possibilities between the involved partners. Some of the mechanisms in WP3 directly evaluate architectures considered in WP4 e.g. the CompSOC platform proposed by TUE relying on Xilinx Microblaze processors as one of the components. As WP3, contrary to WP4, considers also off-chip solutions not every mechanism can be directly transferred between WPs, however many of them can serve as inspiration for the solutions as many of the problems repeats across the domains.

In the next section, we present the overview of the conducted work using as examples four mechanisms developed in the scope of the WP3 which were presented as demonstrators during the review meetings and conferences.

Dynamic sharing of Virtual Channels in Network-On-Chip (TUBS)

TUBS proposed in T3.6 an alternative approach for providing efficient service guarantees in NoCs for mixed-critical real-time systems. It concentrates mainly on the virtualization and isolation layer from Figure 6. Its aspects, however, consider also dynamic reconfiguration and networking. The mechanism combines the global scheduling for the end-to-end guarantees, with the local arbitration performed in routers. TUBS introduces scheduling modules, called resource managers (RMs), with which applications have to negotiate their accesses to interconnect. Synchronization is achieved using control messages and a dedicated protocol. RMs conduct a global, priority based scheduling and grant transmissions access to the NoC for a predefined amount of time. This allows exclusive access to the NoC, hence reducing blocking and decreasing the size of necessary buffers in routers. Moreover, it supports sharing of the same VC between transmissions with different criticality levels while preserving service guarantees. This can be used to decrease the number of required VCs in a system or to

increase the number of hosted applications. The efficiency of the solution derives from work-conserving priority aware scheduling which directly addresses requirements of transmissions with different criticality levels.

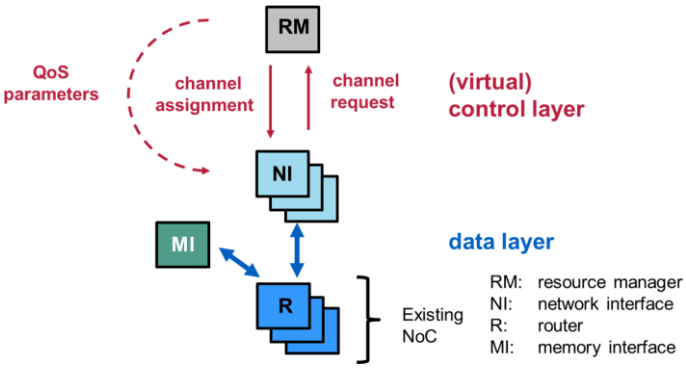

Figure 6: Structure of the SoC system with the global and dynamic arbitration introduced by TUBS

The average utilization is additionally improved through dynamic budgeting. This allows to drastically reduce the average latencies (i.e. temporal over-provisioning) compared to other time-triggered architectures. Therefore, RMs allow to overcome the drawbacks of previously described approaches through reducing hardware overhead compared to non-blocking routers as well as temporal overhead compared to TDM. The introduced solution does not require modification of routers and therefore can be used together with any architecture utilizing non-blocking routers. The demonstration followed on the IDAMC platform, closely related to SoCRocket from WP4. However, the proposed methodology is universal and can be applied to the majority of the NoC-based multi- and many-core systems. Safety of the mechanism must be confirmed with tools for the formal verification of the worst-case behavior e.g. end-to-end latency. Introduced efficiency drives new platform specific solutions which can be accompanied by domain specific tools for mapping, testing and deployment. The demonstrator was presented during the first review meeting in Munich, May 21-22, 2015.

Multicore based communication platform (Infineon)

Infineon introduced in the scope of the task T3.2 mechanisms for the AURIX Microcontroller equipped with Ethernet connectivity. The main goal was to enable AURIX as a safe and efficient platform for mixed criticality communication. The work concentrates mainly on the network layer form Figure 7.

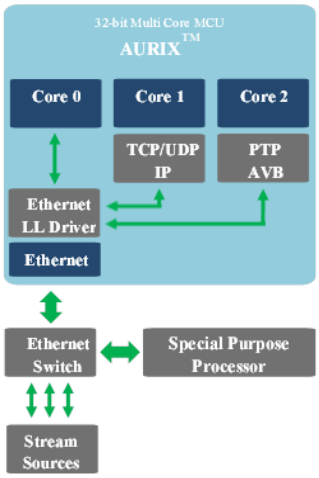

Figure 7: Infineon setup for the safe handling of mixed-critical communication

The first period of work considered the evaluation of the Aurix platform, also in the scope of WP4. The system shown in Fig.3 offers high computing throughput and scalable performance, because it can process several tasks in parallel on different independent cores. Furthermore, AURIX offers a system to protect memory resources from accesses, which are unintended or malicious. As result the system can execute protected communication over several parallel “channels”. These channels shall be separated due to the requirement that they can transfer information of different content and of different behavior profile. So Infineon introduced a solution using methods which support the freedom from interference between communication channels. Such solution was presented in a demonstrator.

Based on the analysis, the following extensions and mechanisms were implemented for RTEs:

- Isolation at application level
- Separation of message flows
- Isolation at stack level to enable deterministic behavior

Isolation at application level is achieved by putting the code images and the related data into different memory areas. Accesses to these areas are controlled and protected by AURIX HW. Separation as method is used in the demonstrator to build a secure environment for the applications in the multicore AURIX. The demonstrator

supports the isolation at two levels: separation based on destination / source of the message flow and separation on the type of messages.

Consequently, the presented demonstrator introduces a setup which is able to isolate the message flows, using the HW protection mechanisms of the AURIX HW. The units build messages which are identical with the frames provided by the Ethernet protocol. The separation of the incoming message flow is the base for all isolation activities between messages. This is achieved using two decision units which must decide which type of messages has been received. In the demonstrator, two different targets can be addressed: TCP/UDP or AVB/PTP packages. The demonstrator was presented during the 1st consortium conference in Vienna, September 28,29 2015.

Safety-critical wireless communication in platooning (NXPNL)

NXPNL in T3.5 proposed a demonstrator considering the safety-critical application of wireless communication based on WiFi-p (IEEE 802.11p) in the platooning of commercial vehicles, see Figure 8 – off-chip networking technology. The solutions concentrates on the newtork and dynamic reconfiguration layers from Figure 5. Business-wise, platooning in the logistics market is forecasted to be an attractive segment in the adoption of automated driving and WiFi-p, whilst providing additional benefits in traffic efficiency, safety and fuel-economy / CO2 reduction. A consortium consisting of DAF, TNO, NXP Semiconductors and Ricardo realized a laterally automated platoon with headway distance of 0.5 sec (11m) @ 80 km/h, on the foundation of the Cohda Wireless MK5 On Board Unit empowered by NXP's RoadLINK dual-tuner chipset.

The custom-development on the MK5-based subsystem provides redundant and low latency platooning-control-information, bi-directional audio between the trucks and video see-through from 1st to 2nd truck, using multiple service channels in the licensed ITS band at 5.9 Ghz. In anticipation of possible future congestion in this band, a number of decentralized congestion control mechanisms as proposed in ETSI standards were investigated, simulated and prototyped. The system was demonstrated on the road in the EU Truck Platooning Challenge 2016, providing ample evidence of real-life performance and suggestions for future work. The demonstrator was also presented during the 2nd review meeting in Gothenburg, June 16-17, 2016.

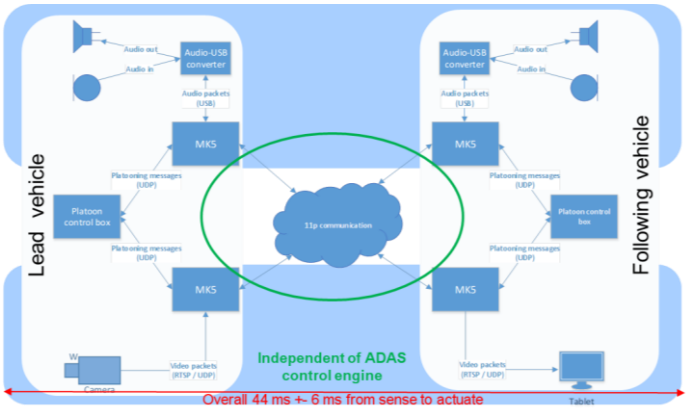

Figure 8: Presented by NXPNL setup with WiFi-p subsystem with mixed criticality data-streams (4 antennas/channels offer redundant messaging, bi-directional audio and video, ETSI standards compliance, diagnostics and security)

NXP/Freescale Embedded Remote Procedure Call (eRPC) Implementation

NXP/Freescale in T3.3 provides the cross-core Remote Procedure Call implementation as the base enablement software for many SoC platforms, including selected SoCs considered in WP4. The Embedded Remote Procedure Call (eRPC) library has been designed and implemented. The RPC is a mechanism used to invoke a software routine on a remote system via a simple local function call. When a remote function is called by the client, the function's parameters and an identifier for the called routine are marshalled (or serialized) into a stream of bytes. This byte stream is transported to the server through a communications channel (RPMsg elaborated in T3.2 task, or TPC/IP, UART, etc). The server unmarshalls the parameters, determines which function was invoked, and calls it. If the function returns a value, it is marshalled and sent back to the client.

RPC implementations typically use combination of a tool (eRPC generator) and IDL (interface definition language) file to generate source code to handle the details of marshalling a function's parameters and building the data stream. The tool also generates code for a server side shim that knows how to unmarshall a request and call the appropriate function. An IDL file is used to tell the generator tool about data types and RPC services.

The NXP/Freescale eRPC infrastructure offers:

- Multi-platforms: erpcgen tool is designed to be executable on most used operating systems (MacOS, Windows® OS, and Linux® OS).
- An easy way to create a client/server application.
- The server application can be blocking (when server serves only for client requests) or non-blocking (when server is executing also another code as the client requests).
- A simple changing transport layer (same applications can use a different transport medium).

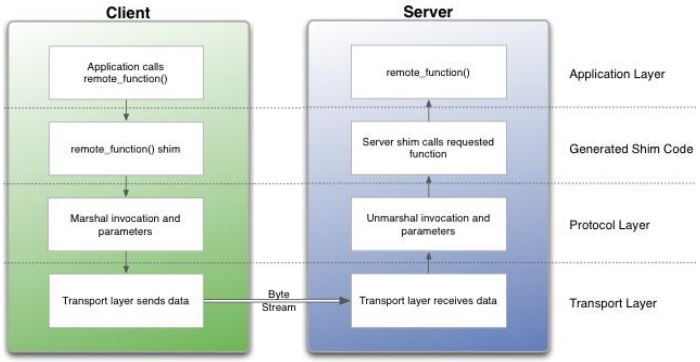

Figure 9: NXP/FSL eRPC Layers

Main eRPC features:

- Lightweight but scalable, targeted to embedded systems
- Small generated code size to allow usage on multicore parts with small memory size
- Abstracted transport interface
- Serialization layer is abstracted and replaceable
- Small size of serialized data
- Designed to work well with C, but flexible enough to support object-oriented languages
- Asynchronous notifications from server to client
- Multithreading of servers when built with an RTOS
- Unique specification of a function to be called
- Provisions for matching response messages to request messages
- Versioning of services
- Minimize any latency impact

The first implementation has been done and available publically at GitHub:
https://github.com/EmbeddedRPC/erpc
A separate GitHub repository has been created to concentrate all guides for running NXP/Freescale multi-core SW (eRPC, RPMsg/RPMsg-Lite) in heterogeneous multicore parts:
https://github.com/EmbeddedRPC/erpc-imx-demos

The demonstrator was presented during the 2nd consortium conference in Paris, September 27-28, 2016, showing the eRPC calls between Linux and FreeRTOS applications using the RPMsg transport layer.

IV. AUTOMOTIVE APPLICATIONS WITH SERVICE ORIENTED ARCHITECTURE (SOA)

Jan van Deventer, LTU

One of the technologies of interest to EMC² is Service Oriented Architecture (SOA) and its integration in multi-core embedded systems used in applications with mixed criticality. This interest is driven by the promise that SOA components do not have to be fixed at design time, can evolve with time and can begin to interact at runtime because they are loosely coupled and late binding. However, issues of security, dependability, interoperability and real time are also crucial to EMC² and are being addressed in the project. Demonstrators are concrete instances of these concepts, two of which we present here.

They both use as a structure for the service oriented architecture the open source Arrowhead Framework [1]. This framework is the product of another Artemis Joint Undertaking project called Arrowhead. It requires only three core services while suggesting several other support services. The three core services are the registry service, the orchestration service and the authorization service. The first one, the service registry, keeps track of the services available. Its service discovery concept is similar to DNS and referred to as DNS-SD. The orchestration service provides the best service provider to the request of a service consumer that is looking for a specific service. The authorization service insures that the

exchange of services is allowed. Together, they form a secure local cloud with low latency (Figure 10) [2].

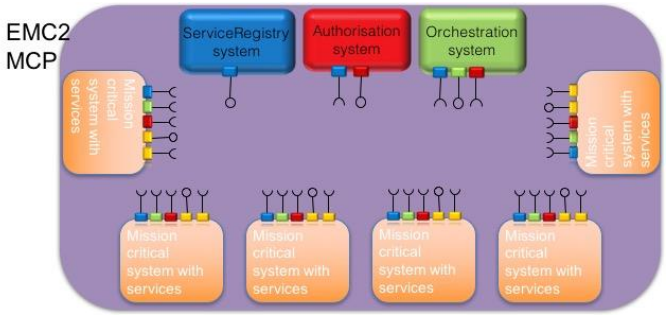

Figure 10: Diagram of an Arrowhead Framework local cloud with its three core services and other application services

VOLVO and LTU have built a climate control demonstrator (Figure 11). A multi-core Xilinx Zynq development board holds the climate control service within an Arrowhead Framework local cloud. The development board communicates via Ethernet to a short-range wireless gateway on to wireless nodes. These nodes are a temperature sensor, an air blower and two ventilation duct flap servos. The services offered by the sensor and actuators are registered and discovered correctly. The demonstrator was shown at the EMC2 second review in Gothenburg and at the Artemis Technology Conference in Madrid of this year.

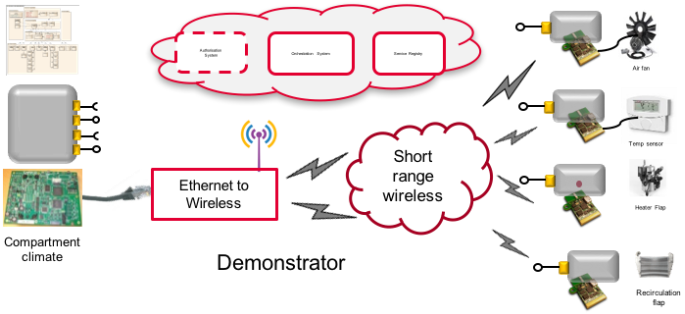

Figure 11: Multi-core ECU for climate control with SOA and wireless sensors and actuators

Students at Luleå University of Technology are developing another demonstrator: a car that has drive-by-wire with active safety (i.e., ABS, electronic stability control (ESC)) and that can offer services, such as vehicle dynamics simulation validation over the Internet (Figure 12). To fit in EMC², it has a main electronic control unit (ECU) with a multi-core embedded system. The car had I/O nodes at each wheel with a standard CAN bus to communicate over to the other nodes and the main ECU. These nodes are simple electronic control units and therefore are referred to as I/O nodes rather than ECUs. The car is, for safety and budget reasons, a 1/5 scale model car (Figure 13).

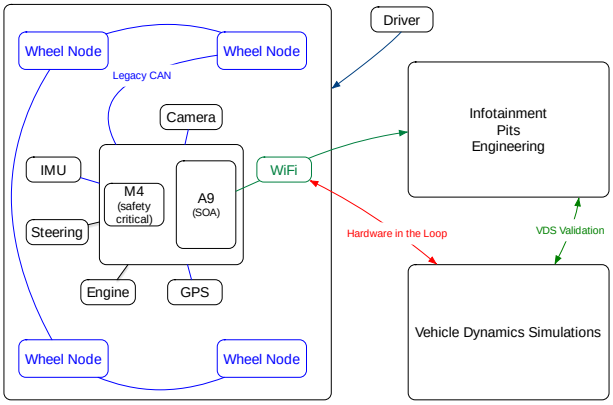

Figure 12: The drive-by-wire with active safety demonstrator concept

Figure 13: 1/5 scale model car with independent brakes

The actual implementation on the model car has a main ECU based on an NXP iMX6s SoloX with a heterogeneous multi-core processor containing a Cortex M4 and a Cortex A9. The I/O nodes are Cortex M0 produced by ST Electronics. The nodes broadcast CAN messages to the M4 and vice versa. The M4 and A9 communicate with each other via a messaging protocol called RPMSG (Remote Processor Messaging). The M0 nodes are interrupt driven and return to an idle state when no specific events need to be handled. They measure wheel speed, actuate the brake calipers via a servo motor. They also make the decision to reduce the brake torque when the wheel tends to lock (distributed ABS). The M4 runs under real time operating system (Free RTOS) and handles the critical aspect of the system. The Cortex A9 runs Linux and offers SOA using the Arrowhead Framework.

Although taking other courses, the ten students were able to design and build the first hardware prototypes and software drivers to prove the concept within seven weeks. Rotating a wheel generates a wheel speed measured by the M0. This data is broadcasted to the M4, which further transmits it to the A9. The A9 offers this information as a service over Ethernet to a PC running the web browser Firefox with its Copper extension. The protocol used in this case is CoAP. The functional implementation was demonstrated to EMC² stakeholders on October 21, 2016. During the following eight weeks, they are to develop the application specific functions such as ABS and ESC. The vehicle states service will be used to validate or reject vehicle dynamics simulations performed in Adams Car using the ISO 19364 standard.

Although the automotive examples illustrate well the concepts of EMC² and SOA, they more importantly provide platforms where the EMC² partners can discuss and argue their own paradigms of Embedded Multi-Core systems for Mixed Criticality applications in dynamic and changeable real-time environments. An example of such discussion is: should the levels of criticality be on different cores or could they be found on the same cores? The model car implementation separates criticality levels on different cores: safety critical aspect on the M4 and non-safety critical aspect on the A9. A variation could move the Inertial Measurement Unit (IMU) needed for ESC from the M4 to the A9 where it would be offered as an IMU SOA service to the M4. From an EMC² point of view, that means that there would be mixed criticality on a single core, which awakes real-time and dependability discussions. Application wise, this refinement would allow the testing of the unmodified active safety software on the target hardware (M0 and M4) in a hardware-in-the loop simulation using the actual vehicle as a service provider.

Nowadays it is easy to find real life examples where cars connected to the Internet have their control taken over by hackers. Since the model car's A9 is connect to the Internet, one should ask if hardware authentication and the Arrowhead Framework authentication are able to protect against that. To enhance this, the Arrowhead Framework has, with its concept of Local Cloud, a Gate Keeper service. When it comes to interoperability, the Orchestration service transparently invokes a protocol translator (e.g., CoAP-HTTP) between service providers and consumers. Dependability is another issue of interest. At the EMC²'s second review in June 2016, Virtual Vehicle demonstrated how dependable SOA (dSOA) could intervene when a sensor malfunctioned and the SOA system took over to correct for that. Their demonstrator was implemented on an Infineon's AURIX microcontroller using the Robotic Operating System (ROS).

EMC² is addressing relevant issues with SOA. Having concrete examples enables the illustration of the concepts as well as to foster discussions about how SOA on multi-core embedded systems with mixed criticality could be implemented.

[1] IoT Automation - Arrowhead Framework, Delsing J. Ed., CRC Press Feb 2017.
[2] Delsing J., Eliasson J., van Deventer J., Derhamy H., Varga P., Enabling IoT Automation Using Local Clouds, 2016 IEEE 3rd World Forum on Internet of Things (WF-IoT), Reston, USA, Dec, 2016.

V. EMC² USE CASE "HIGHLY AUTOMATED DRIVING"

Raul Correal, IXION

This article is about a system architecture that exploits the potential of existing technologies around highly automated driving, which has been developed in task 2 of WP7. Preventing accidents and automating tasks in progressively more complex driving situations requires innovation in perceiving the environment, the vehicle-state, and the reasoning about them. This can be done by the use of multiple sensor technologies such as radar, computer vision, LIDAR or even ultrasound, complementing and reinforcing each other, and communicating over a high bandwidth reliable network. The optimal use of such a huge amount of heterogeneous information requires innovations in both the involved component technologies and in the system architecture concepts. It is then

mandatory to propose tailored solutions that take into consideration the heavy computational load imposed by such systems, and to adaptively combine it with the existing and upcoming platforms to significantly reduce the number of ECUs that would otherwise be required.

Mixed criticality management in automotive embedded systems is one of the hottest topics in the field. Thus, a commercial vehicle include nowadays a platform for time-critical embedded control-units (engine control, ADAS, highly automated vehicles), a platform for entertainment (navigation and multimedia) and even a platform for smart phones communication. In a parallel line, the progress in standalone and cooperative control systems towards highly automated vehicles is becoming a reality in experimental prototypes. It is then necessary to provide adapted dependable solutions that take into consideration the heavy computational payload of such systems and adaptively combine them with the existing platforms.

The main objective of this work is to investigate, implement and evaluate a system architecture that best exploits the potential of existing technologies around highly automated driving. To that end, the EMC2 architecture and tools enable the scheduling of time-critical and less time-critical high-performance functionalities on the same system, tested in real driving scenarios. This work goes beyond the state of the art in communicating ADAS and Highly Automated Vehicles through the design, development and validation of a solution optimizing the HW/SW resources of the multi-core embedded cloud proposed in EMC².

To that end, an urban commuting scenario has been identified, where safety for driver/passengers and predictability for other road users have to be guaranteed. Vehicles travel along the same roads around the same time of the day in a daily basis throughout the year. Often trajectories involve travelling from small roads within the rural and peri-urban area towards peri-urban agglomerations, which implies in many cases both merging to major roads and overtaking other vehicles.

In this use case the commuting vehicle is driven in a manual manner in a first stage. During this process, the embedded software learns the vehicle behavior and simultaneously localizes and builds a navigable map for autonomous navigation. When enough experience is acquired, the vehicle can inform the driver that it is now ready to drive autonomously in a second stage. When driving autonomously under the supervision of the human driver, the vehicle continuously monitors the integrity of the information provided by its sensors. In case of loss of integrity, it informs the driver that confidence is not high enough to continue to drive autonomously. In this case, an adapted HMI would suggest the driver to take-over control of the vehicle. In this context, if either some sort of functional mismatch is detected by the system and the driver prefers not to take-over or the driver simply feels sick, a stop over a safe location is automatically performed. The decision system is in charge of guaranteeing that the man-to-machine transition is secure, in the first case, and of guaranteeing smoothness and safety in the automated vehicle maneuver.

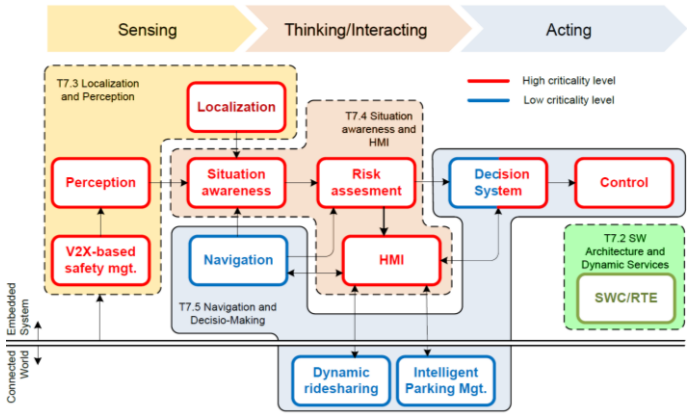

Figure 14: Functional architecture of the highly automated driving urban commuting use case

Vehicles which merge do so by finding the proper slot in the flow of vehicles on the major road. Once merged in the traffic, the automated vehicle reproduces the learned behavior if it is alone on its lane; otherwise it regulates its inter-distance with the vehicle in front of it. Besides, overtaking may be necessary if a vehicle is accidentally stopped in the middle of the lane or if it is just running too slowly for the cruise speed of our vehicle. It has been shown that both automatic merging and overtaking can be facilitated through the use of embedded V2X communications technologies.

The objective of this work has been therefore to use EMC² architectures and tools in real-life tests, using current advanced sensing, navigation and co-operation

functionalities for highly automated vehicles. Both simulation and real-life experiments have been made to show the potential of multi-core new service oriented architectures in the context of highly automated and connected driving. As a result of this, significant contributions with real-life experiments have been made in this still immature but promising field.

The proposed novel architecture of EMC² finds a solution to the fact that the increasing number of ECUs is hardly sustainable. Additionally, a key issue for driver acceptance and quick market introduction of such real time complex controlled systems is to assess its performance under uncertainty. In this connection, new functional safety and integrity requirements are appearing (ISO 26262).

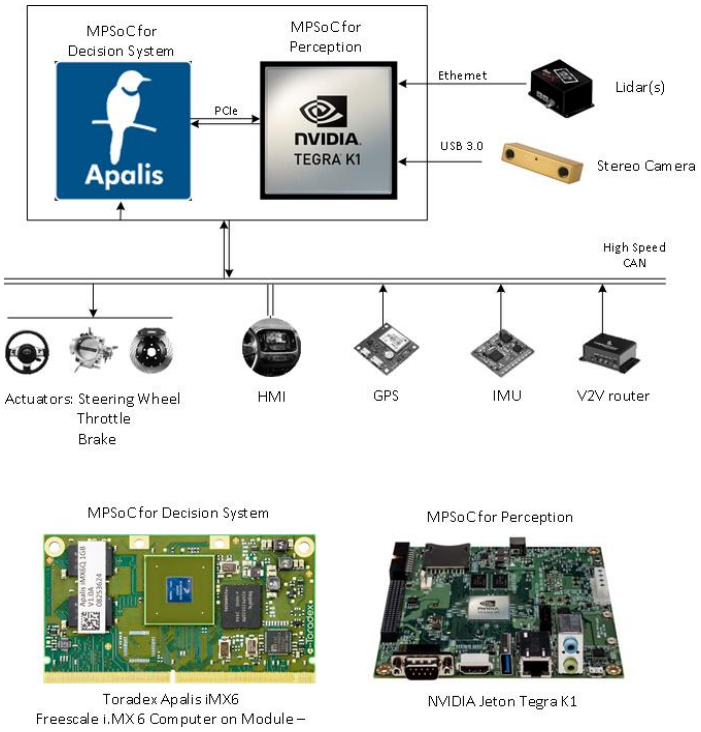

Figure 15: HW architecture of the MPSoC for UC T7.2

In consequence, the main objective of this work is to use EMC² architectures and tools in real-life tests, where a highly integrated computing unit has been developed, using current advanced sensing, navigation, decision making and cooperative functionalities for highly automated vehicles. The resultant design includes a SW architecture enabling the implementation of efficient algorithms in multi-core environments while preserving a high level of safety and reliability.

More specifically, multiple existing systems with different characteristics in terms of computational cost, control cycle period, or level of criticality (ISO 26262, IEC 61508) have been integrated. As a result, (i) an increased integration along with a (ii) consistent separation are natural requirements for the resulting mixed criticality runtime environments.

- **Increased integration**: Highly intensive computational components such as perception, map-based localization, decision making and V2X communications have been integrated in a reduced set of computing platforms, using a hybrid combination of different technologies within the same HW platform. The existing set of verification & validation artefacts for isolated components have been exploited for the resulting configuration in order to reduce V&V costs and efforts.
- **Separation**: One of the objectives for a mixed-criticality solution is to provide sufficient separation between elements of differing criticality levels to guarantee that the lower criticality elements cannot interfere with the functioning of the higher criticality elements. In particular, navigation related SW components have a lower criticality than most of the embedded subsystems. However, there are some specific components that combine high and low critical systems (namely, the decision system, which includes a time-critical motion planner, and a learning process that can be handled with a much lower criticality). System fault tolerance and resources requirements at local level have been taken into consideration with modelling approaches.

VI. EMC² USE CASE: "DESIGN AND VALIDATION OF NEXT GENERATION HYBRID POWERTRAIN"

Georg Macher, AVL

In the automotive domain current premium conventional and electric vehicles are characterized by several tens of distributed control units (more than 90 electronic control units with close to 1 Gigabyte software code [1]), with a complexity level of the E/E architecture significantly constraining vehicle performance improvements. Nevertheless, these systems are responsible for 25% of vehicle costs and an added value between 40% and 75% [2].

To achieve further benefits a further enhancement of the control functions and re-design of the current E/E architectures will soon be required. High-performance domain controller based on multi-core systems enable the deployment of more advanced control strategies providing additional benefits for the customer and environment, but at the same time the higher degree of integration and criticality of these control application raise new challenges. These factors cause multiple cross-domain collaborations and interactions in the face of the challenge to master the increased complexity and ensure consistency of the development along the entire product life cycle.

Consequently, UC7.3 focuses on enhancement and adaptation of vehicle electric / electronic architecture to better suit (a) the requirements coming from new applications fields such as higher degree of electrification and hybridization driving, and (b) opportunities coming from SOA approaches. UC7.3 provides an integration platform for technologies developed in the other WPs. The generic automotive E/E board net platform enables (a) the integration of new architecture concepts such as Service-Oriented Architecture (SOA), and (b) provides an integration platform to evaluate the outcomes from the different other WPs in a realistic environment.

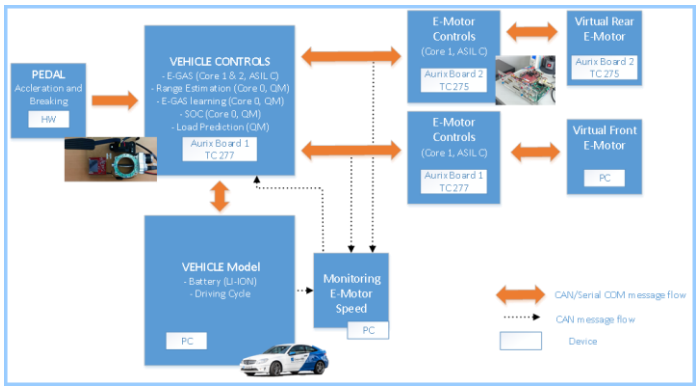

Figure 16: Generic automotive E/E board net demonstrator

The generic automotive board net demonstrator provides the basis for independent and mixed-critical SW application integration on multi-core computing platform and integration of 4 independent SW stacks from different SW supplies, as demonstrated at 2nd year review meeting.

Beside the generic E/E architecture demonstrator, UC7.3 also focuses on dependability aspects, safety and security. An integrated dependability framework for providing product and process based safety and security argumentation, as well as, a contract-based approach for conditional safety certificates (M2C2 demonstrator) are integral part of UC7.3.

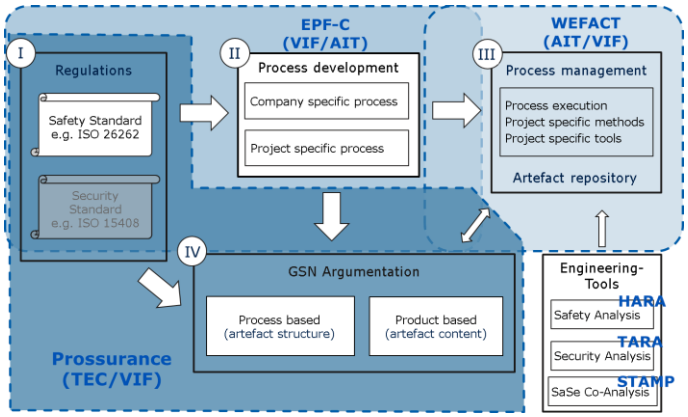

Figure 17: Integrated dependability framework.

This framework provides a comprehensive framework for safety & cybersecurity co-engineering, especially addressing the following aspects: (a) safety & cybersecurity assurance cases with the capability to generate and efficiently manage dependability information along the development lifecycle and (b) new safety and security concepts and their respective implementations. Thus, the

framework shall help to ensure an efficient use of the components and E/E infrastructure and finally develop advanced and dependable control strategies for the end user.

Nevertheless, future needs of vehicle development rise the need for service-oriented architectures and dynamic reconfigurations. The dynamic compositions and reconfigurations (due to web-based functionalities and update-over-the-air functions) of these systems hamper the use of established engineering approaches, which cannot be applied without further ado. Current approaches assume, that a system as well as its usage context is completely known and can thus be analyzed thoroughly at development time, which does no longer hold true.
As a result, a shift of parts of the safety certification activities to runtime, where the complete information can be obtained and uncertainty can be truly resolved, is focused for the remaining project duration. The idea behind M2C2 contracts, a safety certification, which are created at development time, but evaluated at runtime by the systems themselves will be integrated in the UC7.3 demonstrator. This will merge all activities of UC7.3 for the final project review.

The efforts of UC7.3 allowed supporting a generic automotive board net demonstrator, which enables the coaching of SW developers (industrial engineers and scientists) for multi-core constraints, mixed-critical SW application integration and different SW supplier alignment. Furthermore, the integrated dependability framework invests ensure to further coalesce SW and safety engineering and strengthen the integration of these engineering domains.

[1] Christof Ebert and Capers Jones. Embedded Software: Facts, Figures, and Future. IEEE Computer Society, 0018-9162/09:42-52, 2009.
[2] Giorgio Scuro. Automotive industry Innovation driven by electronics. http://embedded-computing.com/articles/automotive-industry-innovationdriven-electronics/, 2012.

VII. COTS Multicores in Avionic Applications


Sascha Uhrig, Airbus


Johannes Freitag, Airbus


Gerhard Fohler (TUKL)


Ankit Agrawal (TUKL)

Emerging safety-critical hard real-time systems like UAVs or autonomous driving applications require integration of many new functionalities. For example, in the future, autonomous helicopters (and autonomous aerial vehicles like air ambulances) targeting the urban city environment need to fly at low altitudes especially when picking-up or dropping-off either people or shipments. Implementation of such emerging applications requires even more performance from avionics hardware, without increasing or even better decreasing size, weight and power (SWaP). These requirements cannot be fulfilled by existing avionics hardware that use current avionics single-core processors.

COTS multicores can provide the requested limitations on SWaP as well as the requested performance, but introduce resource contentions on co-executing tasks due to sharing of hardware resources like on-chip network, memory sub-system etc. These shared resources result in timing dependencies amongst tasks leading to non-determinism of their execution time, i.e. tasks may not meet their deadlines. This behaviour can cause a serious safety issue on avionics systems. The CAST-32 position paper [1], jointly from EASA and FAA, further highlights the severity of the problem and as a first step, proposes to limit the number of active cores to just two.

A naive solution is to use COTS multicore processors with only a single active core and keep all other cores either power down or idle. This results in waste of resources, while also defeating the purpose of increasing performance with the same SWaP. Solving the contention problem on COTS multicores requires an integrated approach combining contention-aware scheduling of applications and mechanisms that ensure bounded interference in COTS multicores.

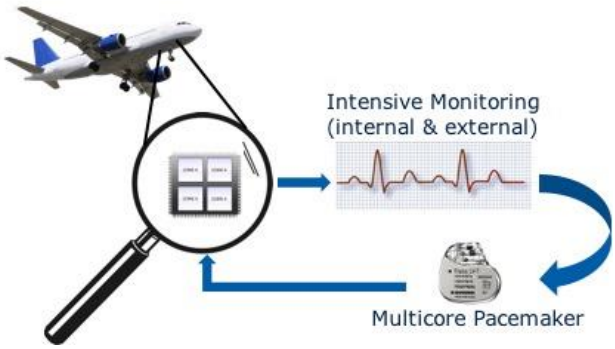

Figure 18: Overview of the methods developed in WP8.2

In WP8.2 of the EMC² project, we developed two methods adhering to the integrated approach, a runtime monitoring method and a static/dynamic scheduling method. For the runtime monitoring, a single critical core of a COTS multicore sends continuously data of the task's progress to an external safety-net processor (possibly DAL-A certified). The external processor comprises a database of performance fingerprints. Based on the executing task and the performance fingerprints it decides if the executing task is on time and will meet its deadline. If a delay is recognized one or more of the non-critical cores will be temporarily suspended or at least slowed down. This gives the critical core the possibility to speed-up execution and to keep its deadline anyway.

The second method allows using all cores in a COTS multicore as critical cores. For each fixed time interval called slot, each core executes tasks based on a schedule table. In each slot, if any task exceeds its assigned time budget and memory access budget, it may result in other co-executing tasks missing their deadlines. Therefore, the core-level scheduler suspends tasks that exhaust their budget reservations, until the start of the next slot. This case of possible runtime suspension is already accounted for in the offline scheduling phase, thereby enabling concurrent execution of critical tasks on more than one core.

Currently, no mixed-criticality (DAL A-E) multicore application in avionics is known, as no COTS multicore has been certified for multi-application use. The black-box nature of COTS multicore architectures combined with the lack of usage data impedes the use of COTS multicores in safety-critical avionics (e.g. DAL A). One way to overcome this issue is to use COTS multicores for lower-criticality applications like DAL C-E, and collect in-flight experiences. This data can then aid in certification of COTS multicores in avionics for DAL A/B applications [2]. The static/dynamic scheduling and the monitoring methods, can aid in this direction as they complement each other and it is possible to combine them. The monitoring method can be coupled with an underlying static/dynamic scheduling and act as an additional layer to ensure that the COTS multicore is executing tasks as intended. In case of any abnormality, the external processor can either take mitigating actions or simply inform the other nodes in the system and log the incident.

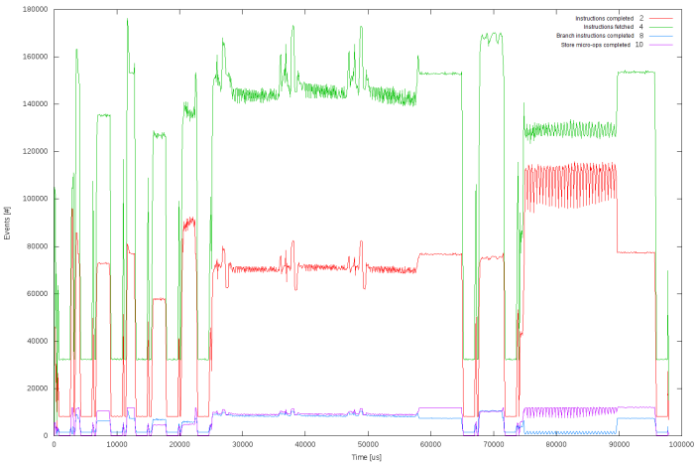

Figure 19: Example fingerprint of a benchmark application

The monitoring method is implemented for a DAL-C application inside an avionics computer. Inside the FPGA a safety net processor is tracking the application executed on the dual-core NXP P5020 processor. An example fingerprint of a benchmark application is shown in Figure 19. The different characteristics of the application's progress can be identified easily. In summary, the presented methods have the potential to pave way for use of COTS multicores in safety-critical domains like

avionics, such that tasks are executed deterministically considering resource contentions.

[1] CAST-32 Multi-core Processors. FAA's Certification Authorities Software Team, May 2014.
[2] R. Fuchsen, "How to address certification for multi-core based IMA platforms: Current status and potential solutions," 2016.

VIII. FPGA PARTIAL RECONFIGURATION IN SPACE APPLICATIONS

Elena Terradillos, Tecnalia

Manuel Sanchez TASE

Within the MPSoC Hardware for Space use case in EMC² WP9, TECNALIA and TASE work together to provide a solution so that partial reconfiguration of programmable devices can be done in a reliable way in the space environment. The target programmable devices are SRAM-based FPGAs and, in particular, the Xilinx Virtex-5QV FPGA which is the highest density and performance space-grade FPGA in the market built with radiation-hardened by design technology to provide intrinsic hardness from SEU (Single Event Upset) and SET (Single Event Transient) to select elements of the device.

SRAM-based FPGAs have become increasingly attractive for space-based computing platforms due to: (1) their on-orbit re-configurability which extends the useful lifetime of the system; (2) their low development cost comparing to ASICs; and (3) they are well-suited to digital signal processing tasks, which are very common in satellite applications. Unfortunately, these devices are susceptible to SEUs which can make fault-tolerant implementations challenging, but not impossible. The Xilinx Virtex-5QV device was rolled out in 2011 and there are many NASA missions with Virtex-5QV programmed to be launched on 2017 and beyond.

If FPGA technology provides the flexibility of modifying a design by re-programming without going through re-fabrication, Partial Reconfiguration (PR) takes this flexibility one step further. PR is the ability to dynamically modify blocks of logic while the remaining logic continues to operate without interruption. It allows designers to change functionality on the fly, eliminating the need to fully reconfigure and re-establish links. This technique presents several advantages: (1) reduce the size of the FPGA device required to implement a given function, with consequent reductions in cost and power consumption; (2) provide flexibility in the choices of algorithms or protocols available to an application; (3) enable new techniques in design security; (4) improve FPGA fault-tolerance; and (5) accelerate configurable computing.

When designing an FPGA-based platform with partial reconfiguration for mixed-criticality systems, there are three elements or key points to be included: a) a Dynamic Reconfiguration Manager (DRM) in the FPGA static partition; b) fault-tolerance elements and design techniques; and c) an external agent to control and monitor the DRM.

A Reliable and Self-Healing Dynamic Reconfiguration Manager is the proposed solution so that partial reconfiguration of a Xilinx Virtex-5QV FPGA can be done in a safe (reliable) way in the space environment. This DRM consists in a MicroBlaze embedded system implemented in the static (that is, not reconfigurable) area of the Virtex-5 device with capability to manage the reprogramming process of several reconfigurable partitions. The DRM is a co-processor of the full system. One of its main elements is the controller to access the FPGA configuration memory at run time (HW ICAP). This capability to access the FPGA configuration memory is used for a double purpose: first, to perform dynamic reconfiguration of the reconfigurable partitions; second, to perform configuration memory scrubbing in order to prevent SEU accumulation in the FPGA configuration memory.

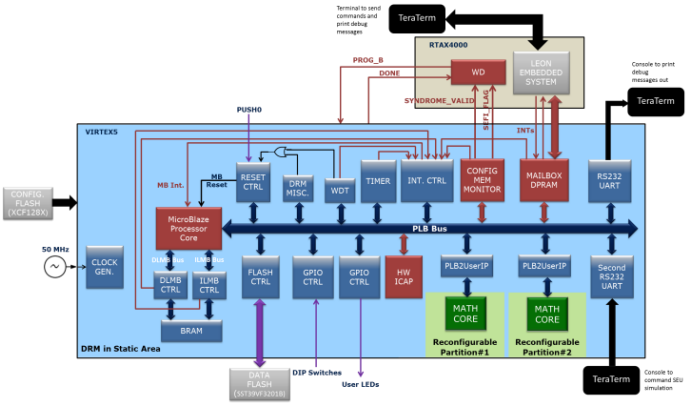

Figure 20: Reliable and Self-Healing DRM architecture diagram

Fault-tolerance elements and design techniques are included so that the DRM can autonomously recover from radiation induced errors. These elements and techniques are:

- FPGA configuration memory management. SRAM-based FPGAs are sensitive to SEUs affecting the configuration memory. This is especially challenging as these upsets may change the behaviour of the FPGA design. Taking actions to mitigate the effects of SEUs in the configuration memory is essential for reliable SRAM-based FPGAs in the space environment and critical systems in general. This task is performed by the MicroBlaze microprocessor with the support of the Configuration Memory Monitor peripheral.
- SEFI detection and mitigation. Single-Event Functional Interrupts (SEFIs) are SEUs that result in the interference of the normal operation of a complex digital circuit. For example: an SEU in the device control circuitry may place the device into a test mode, halt, or undefined state; or the so-called sequence loss in processors resulting, for instance, of an SEU in the program counter leading to an infinite loop. In general, SEFI is a broad term used to indicate a failure in a support circuit within the device and not a failure of the user design. In such cases, a reset of the application or a power off/on cycle is required to recover the full functionality of the system. The external monitor agent forces full Virtex-5 reprogramming when SEFI conditions are detected.
- MicroBlaze soft-core protection. It consists in enabling the MicroBlaze fault-tolerance features to detect errors in the processor memory and perform the corresponding correction actions.
- Partial reconfiguration data integrity check. The purpose of this technique is to avoid loading a corrupted partial bitstream.
- The EDAC (Error Detection And Correction) circuit of the RAM blocks is enabled in the whole design, both the embedded system in the static partition and the designs in the reconfigurable partitions. The SET filters of the CLB flip-flops in the Virtex-5QV FPGA will also be enabled.

The DRM has been designed to be implemented in the Virtex-5 device of the LADAP (L3SoC and Data Processing) platform developed by TASE. The two main elements of this hardware platform are: the already mentioned Xilinx Virtex-5 device, and a Microsemi ProASIC FPGA that will include the main processor of the system.

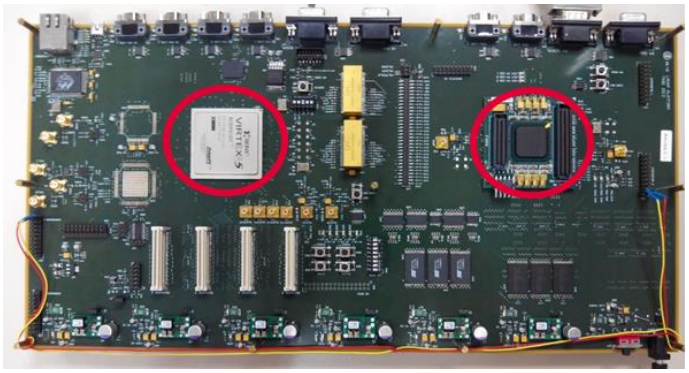

Figure 21: LADAP platform

The DRM is externally controlled and monitored by the LEON processor running in the ProASIC device. The communication interface between both processors has been devised as a mailbox implemented in a shared dual port memory in the Virtex-5 FPGA. The LEON processor is always the communication master. It periodically checks that the DRM is alive and commands the partial reconfiguration actions. Together with a watchdog module, it reprograms the Virtex-5 FPGA when SEFI conditions are detected.

During the EMC² second year review, a demonstration video showing partial reconfiguration working in the Virtex-5 FPGA under command of LEON processor was presented. The fault-tolerance elements have been implemented afterwards and their functional validation is in progress. Autonomous fault recovery of the system will be demonstrated in the final review.

IX. EMC² "INTERNET OF THINGS" USE CASES

Yudani Riobó, Quobis

Internet of Things (IoT) defines a scenario where everyday physical objects are connected to the Internet and interoperate with other objects and systems. The challenge relies on the higher computing resources needed to process the increased amount of data available, as well as safety and security issues in open and dynamic environments. The Internet of Thing living lab focuses in different scenarios to develop new functionalities and improvements. It includes a number of use cases that are providing some interesting results.

Deterministic network of cameras

One of the use cases aims to showcase the potential of open deterministic networking by connecting a variety of local embedded systems to other embedded systems. Open deterministic networks are an enabling technology for future, potentially heterogeneous and mixed-criticality, Systems-of-systems (SoS). There, distributed components or systems will be connected and integrated towards larger systems, while their individual requirements, e.g. regarding safety parameters, are preserved by the novel networking infrastructures.

The first prototype demonstrates a novel approach for monitoring people on airports and measuring passenger data by interconnecting cameras, server and client in a deterministic communication network.

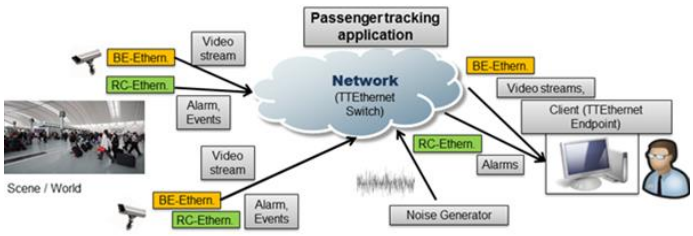

Figure 22: Measuring waiting time by utilising a deterministic network of cameras, server and client

The person tracking application and the additional middleware, both developed in living lab WP12, is used to implement server side the demonstrator's passenger monitoring. Cameras and client are specifically connected via TTEthernet for streaming video data and server side generated events to the client.

The demonstrator provides two different modes of communication namely rate-constrained (RC) and best-effort (BE) communication. The event data including specific alarms are transmitted via the rate-constrained mode as this communication guarantees a maximum delay of one Ethernet frame or packet, if there is one RC link out the specific port. If many RC links are to be sent out the same port (client), then the frame might be delayed longer (first come, first served). A delay of a few Ethernet frames is more than acceptable for the proposed application, so that the third time-triggered mode that guarantees to transmit data without almost any delay is not needed. The server is periodically transmitting the event data to the client. It contains the current on average waiting time monitored by each camera and eventual alarm messages for example when on average waiting time exceeds a user defined threshold. In this demonstrator, waiting time is the time a passenger is correctly tracked by the tracker application.

A noise generator has been connected to the network as well in order to simulate network overload. In such situations the demonstrator shows that event data in rate-constrained mode arrives correctly at the client whereas video directly streamed to the client under best-effort mode may show frame drops. Experiments verifies that deterministic network communication allows safe processing of video and safe event delivery at the client side which is a necessary prerequisite when putting cameras into the control loop of the gate process.

Synchronized low-latency deterministic networks

This use case aims to increase the dependability on communication networks for Smart Grid, focusing on availability, safety, security and reliability. Control, monitoring and critical data rely on these features and thus, on synchronization. Timing data becomes crucial since it is required to provide properly the intercom-

nection between different grid elements and guarantee their functionalities. Next Smart Grid network generation lie on high-accurate, reliable, available and scalable technologies to distribute time. White Rabbit (WR) is intended to be the next-generation deterministic network based on synchronous Ethernet, allowing low-latency deterministic packet routing and high precision for timing transmission. WR networks are composed by a master node that provides the main time and frequency reference, fibre switches, twisted-pair copper, and slave nodes.

WR allows a very precise time-tag data measurement and data triggering acquisition in large installations at the same time that data are also transmitted through the same network. The demonstrator of this use case addresses the development of a novel technology for Smart Grid control. In order to validate the design of this demonstrator two prototypes have been developed. The first prototype focuses on evaluating the scalability, synchronization accuracy and the utilization of different time distribution mechanisms, such as WR-PTP and IRIG-B. The second one focuses in assessing the robustness and availability of different services focusing on time distribution.

Accurate & heterogeneous time transfer system prototype

This prototype is composed of a WR-ZEN configured as a time provider connected to a GPS. The WR-LEN boards are connected in a Daisy Chain configuration with a modified White Rabbit PTP Core (WRPC) containing two endpoints.

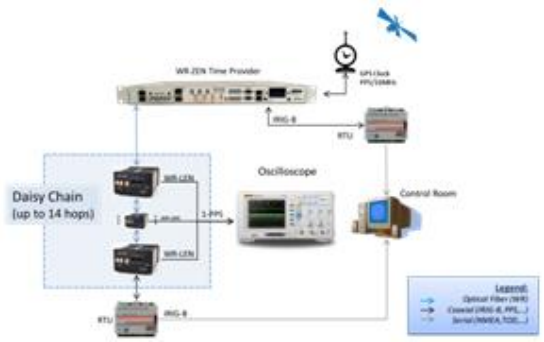

Figure 23: Scalable heterogeneous time transfer system using WR-PTP and IRIG-B

This prototype integrates the project results regarding the determinism of the network architecture and the ones focusing on enhanced QoS definition for critical data packets as timing ones. This demonstrator illustrates the benefits and capabilities of the presented approach for Smart Grid. Benefits of this synchronized, low latency and deterministic networks as enabling technology for new applications or for improving development of existing ones will be shown.

High-availability and fault tolerance prototype

This prototype consists of a WR-ZEN configured as the Grand Master time provider, connected to a GPS, which disseminates time and frequency with 1ns accuracy to a HSR ring network composed of three WR-HSR-Switches. These switches are also synchronized to the WR-ZEN providing a 1-PPS signal perfectly synchronized to the GPS clock 1-PPS.

Time frames are propagated from the WR-ZEN to the first WR-HSR-Switch, where it will insert a HSR tag, duplicate and send both the frames through the ring to the other two WR-HSR-Switches. Each of the switches will receive two copies of the same time frame, making one of them the time reference, and the other one is used a backup in case of failure of the other ring path.

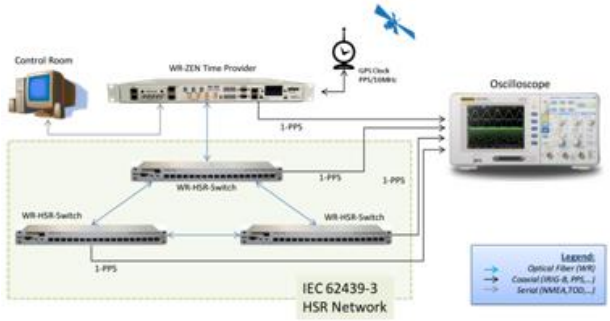

Figure 24: Time transfer fault tolerant system in HSR network

When one of the nodes or path is down, one of the time references is lost, and the switches shall switchover to the previously backup time reference. This is done with approximately zero-time recovery, feature required by Smart Grid high-demanding applications where timing is considered critical.

X. EMC² USE CASE: "VIDEO SURVEILLANCE FOR CRITICAL INFRASTRUCTURE"

Axel Weißenfeld, AIT

Passenger Tracking - Architecture for Dynamic Allocation of Computer Vision Tasks

This article is about an architecture for dynamic allocation of computer vision tasks, which has been developed in task 2 of WP12 and was presented at the International Conference on Distributed Smart Cameras 2016 (ICDSC). The use of reconfigurable computer vision architecture for real time image processing tasks is an important and challenging application in cyber-physical systems with limited resources. This architecture shall be used in WP 11.2 to realize a networked smart vision system using the TTEthernet technology from partner TTTech. This network supports mixed-criticality communication, thereby ensuring that high criticality communication will be guaranteed thereby using single communication networks. Thus, also low-criticality messages are sent over the same network.

Heterogeneous reconfigurable vision processing architectures are attractive and have recently become important. For applications such as passenger tracking at airports (Figure 25), unforeseen dynamic changes of the environment are typical and pose fundamental challenges to computing systems *with limited hardware and software resources*. Hence, vision processing architectures are needed that ensure a certain Quality-of-Service (QoS), maintaining a high level of resource utilisation while achieving the expected system performance. Graceful degradation of system performance ensures to maintain *predictable performance* even with limited resources. Predictable performance in this context means that system performance declines *gradually* as the available resources decrease.

For passenger monitoring, huge amounts of image data need to be processed and interpreted in a network of cameras and computers with rather limited bandwidth and computational resources. In order to cope with this challenge, it is possible to either reduce total image data transfer and throughput or to reduce algorithm complexity and software runtime for image interpretation. We develop a novel reconfigurable software framework that is able to vary both - size of data and computer vision algorithm complexity - in terms of achieving QoS as well as a graceful degradation of system performance when resources become too limited.

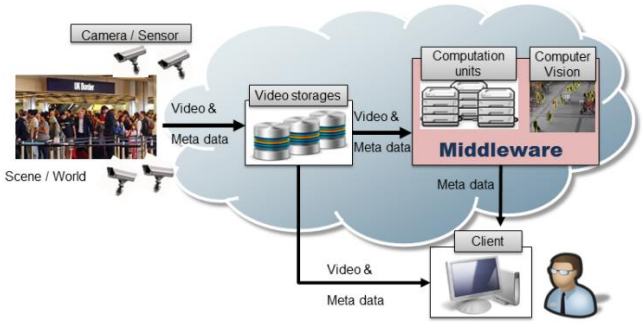

Figure 25: Airport surveillance system with computing units extended by the proposed middleware. IP camera videos are streamed and centrally recorded as well as processed. Videos and Meta data are monitored by the operator on demand.

Proposed Computer Vision Architecture

The proposed computer vision architecture consists of multiple components embedded into individual modules that follow the micro-services architectural style, where each module acts as an autonomous web service communicating via a Representational State Transfer (REST) interface. The framework puts emphasis on distributed processing, so that vision tasks can be distributed over various units. In this way, each module can be configured more easily to sustain a defined QoS and assign the task to another computational node if necessary.

Exemplarily, a technical implementation of the passenger tracking based on the novel computer vision architecture is depicted in Figure 26. A basic tracker setup for passenger tracking might consist of a module chain comprising only three modules: a module providing the video data followed by an object detector algorithm and finally a tracking algorithm. Each module (computer vision task) can be executed on a different processing unit in order to support near real-time.

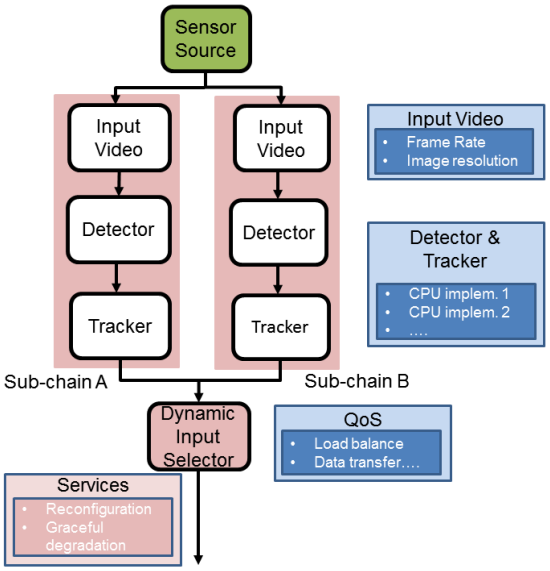

Figure 26: Technical implementation of a passenger tracking application based on the proposed computer vision architecture. The Dynamic Input Selector (DIS) activates either sub-chain A or B depending on the current QoS

The key module which manages the reconfiguration and triggering of the execution is denoted as dynamic input selector (DIS). This module consists of a control and QoS unit, which can be configured by the user. It is a platform-independent solution for supporting chains of computer vision algorithms and reconfiguration in service-oriented soft real-time environments. The control unit triggers a reconfiguration whenever there is a resource variation with impact, which implies to select another sub-chain. These sub-chains are pre-configured, so that a system reconfiguration is achieved without significant delay. Based on quality of service parameters provided by the modules, such as the processing duration per frame, the dynamic input selector chooses to operate between multiple sub-chains. The decision is made by analysing the QoS parameters provided by the preceding modules.

XI. Developing Safety & Security Co-Engineering Methods for Cyber-Physical Systems – Windows of Opportunity for Standardization

Erwin Schoitsch, AIT

Christoph Schmittner, AIT

Zhendong Ma, AIT

Thomas Gruber, AIT

Within EMC² WP6, the AIT is developing a combined approach to safety & security engineering for Cyber-physical Systems (CPS). Since CPS increasingly rely on ICT, attacks from the cyberspace can cause devastating impacts in the physical world. Mixed-criticality and multicore systems present additional challenges for the development of dependable systems. We aim at a holistic approach to safety and security engineering and the development of a security aware safety lifecycle. As a first step we extended Failure Mode and Effect Analysis (FMEA), a well established safety and reliability analysis method, with vulnerability and threat analysis in order to enable a combined consideration of security and safety.

The Failure Mode, Vulnerability and Effect Analysis (FMVEA) method is applied to safety and security analysis of intelligent and cooperative vehicles in the automotive domain, to railway applications (particularly urban transport), and smart manufacturing (Industry 4.0), for the identification of attack possibilities and failure scenarios [1]. The method is compared with Combined Harm Assessment of Safety and Security for Information Systems (CHASSIS) in a case study in automotive cyber-physical systems, with respect to applicability and future research needs [2]. Based on the results from the analysis in the automotive domain, standardization challenges for safety and security have been identified [3]. Besides, we

cooperate with researchers at the Advanced Digital Science Center (ADSC), an affiliate of the University of Illinois in Singapore, to apply safety & security analysis to urban railway systems. In addition we are developing a tool for a partially automated model-based FMVEA.

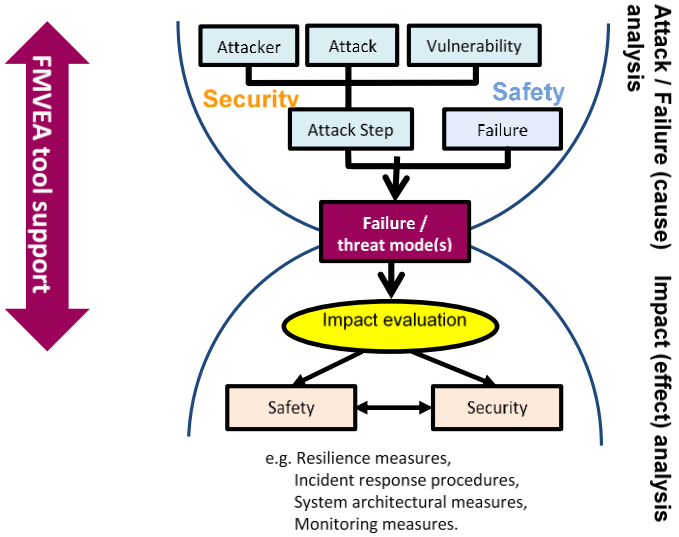

Figure 27: Tool concept for partially automated model-based FMVEA

To manage the workflow of a qualification and certification process, a tool WEFACT (Workflow Engine for Analysis, Certification and Test) maps the requirements of relevant standards and domain- resp. application specific requirements to the V&V and qualification/ certification process. This tool was already used in several use cases of EMC² together with partners like AVL and ViF. Originally only designed for safety analysis, V&V and qualification/certification processes, it is extended to support the complete engineering process and match system security requirements as well.

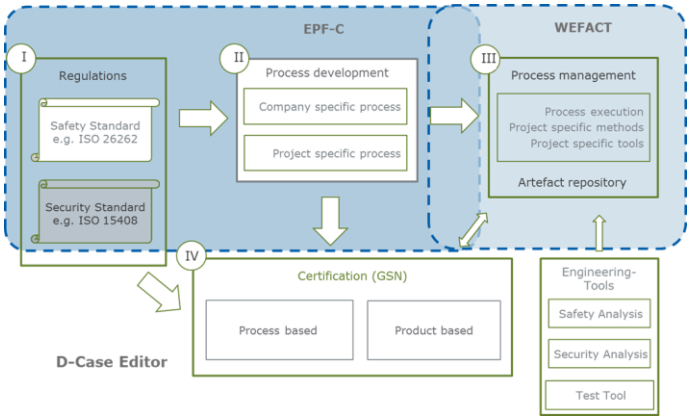

Figure 28: WEFACT tool

The tool allows

- Automated execution of safety and security processes modelled in Eclipse Process Framework (EPF) composer
- Traceability from claims to requirements to evidence
- Integration / interaction of D-Case Editor with WEFACT
- Building a Dependability Assurance Case

<u>WEFACT current operational environment: processes modelled in EPF-C, D-Case Editor, Analysis and Tools</u>

WEFACT allows integrating other (external) tools following the IOS (Interoperability Specification) concept of the ARTEMIS High-Rel Cluster projects. There is a close link to CP-SETIS "Towards Cyber-Physical Systems Engineering Tools Interoperability Standardization", an ARTEMIS-IA driven support action-type Innovation Action in Horizon 2020, trying to harmonize and create a sustainable structure to maintain the different developments towards the IOS for tool interoperability.

<u>WEFACT Framework: Sources of requirements, V&V, external tools, output/feedback</u>

Concerning mitigation of cybersecurity attacks on complex systems and systems of systems, a method called "Rotating virtual machines" was developed. In short, it uses redundancy of virtual machines (VM) to avoid propagation of the consequences of attacks on the system. It is an approach for safety & security co-design of safety-critical multi-core systems to realize security and fault tolerance. Multiple virtual machines (VMs) are switching between active, stand-by and cleansing roles (see Figure 29):

- **Active** VM - there is always one VM providing services
- **Standby** VMs - there are several VMs, in which one of them is ready to become the next Active VM
- **Cleansing** VM - the VM that steps down from the Active state, it is then restored to a clean state to remove any potential malware infection

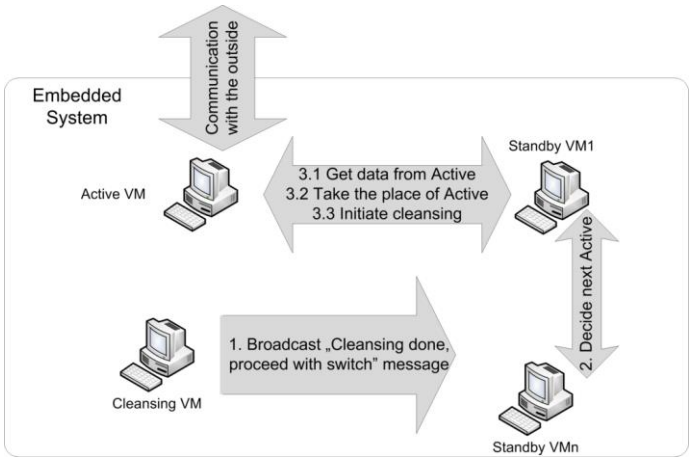

Figure 29: Multiple Virtual Machines to maintain secure system status

The methods and techniques described above (developed in context of task T6.2) are addressing Safety Assurance by Design. However, complex systems and systems of systems have to be open, adaptable systems, a topic addressed particularly by EMC². Therefore, the dependability properties have to be maintained over lifetime, during operation, maintenance and enhancement, as well. This is addressed by task T6.3 "Run-time Certification", lead by Fraunhofer IESE, Kaiserslautern, and University of Kaiserslautern. This is a contract-based approach; the contracts ("conserts", modular conditional certificates) are proven in advance, and when system constituents are added/integrated during runtime, the conserts are checked during run-time if the safety & security requirements will still hold.

This means:

- Parts of the safety certification activities are shifted into runtime, and automated safety checks for dynamic systems integration and adaptation are performed
- The process is based on "ConSerts": modular conditional certificates that are issued at development time for each system
- Support for security assurance between safety critical open systems is investigated as an important part of the research work

This method is particularly useful to be able to check the feasibility of planned system updates caused by security issues (we know. Safety means "never change a certified system if possible", security on the other hand requires regular security updates), so this is a new approach to cope with these diverse system requirements for safety and security.

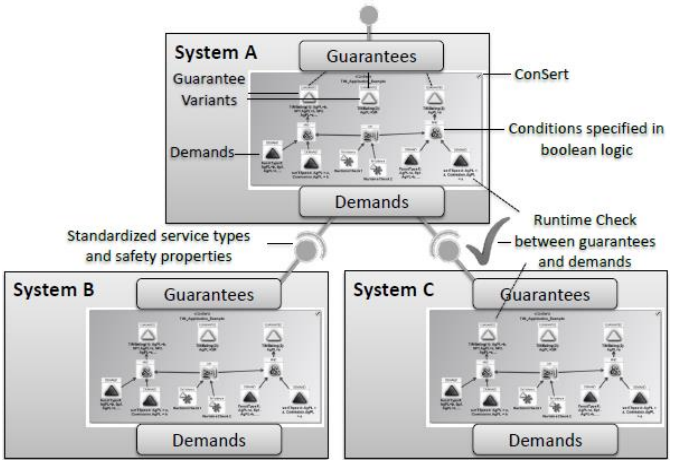

Figure 30: Concerts: Guarantees and Demands

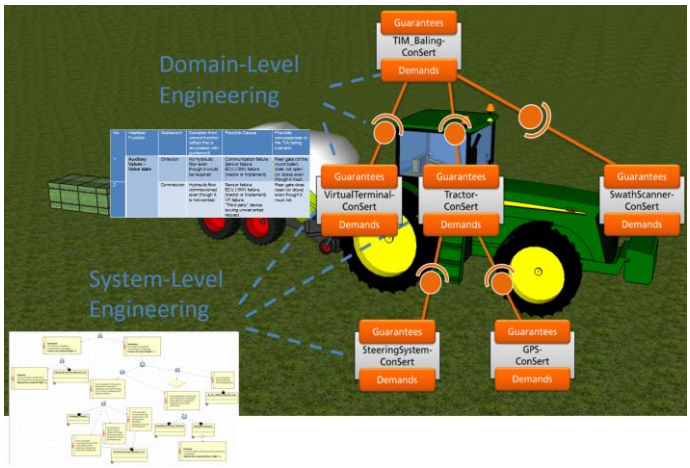

Figure 31: Example for run-time qualification/certification (FhG IESE ©); A tractor with various implements from different providers – do safety requirements hold?

EMC² Impact on Standardization – the currently evolving standardization Landscape
(Erwin Schoitsch, AIT)

On international standardization level, awareness is raising for the topic. IEC TC65 [6], Industrial-process measurement, Control and Automation, had started an ad-hoc group AHG1 to investigate the issue of coordination of safety, security, and was looking at a broad variety of domains and standardization groups starting to think about including (cyber-)security aware safety considerations. This was achieved already partially in IEC 61508, Ed. 2, by a group where AIT was strongly engaged (E. Schoitsch), and is now starting in the railway sector (DKE

in Germany, integrating requirements from IEC 62443 in the railway standards (proposal, addressing EN 50129 and EN 50159 issues) in DIN VDE V 0831-104 "Electric signalling systems for railways – Part 104: IT Security Guideline based on IEC 62443" and in IEC TC44 (safety of machinery, electro-technical aspects). In the meantime, AHG1 has completed its work with a report recommending to prepare an IEC TS on the topic "Framework to bridge the requirements for Safety and Security" and started a new working group IEC TC65 WG 20 under this title. There have been already a few Face-to-Face meetings (one in Vienna at AIT) and work is done via web and telephone conferences (almost monthly), our goal is to keep our EMC² triggered intention to foster safety & security co-engineering and remain on a level to produce a basic safety & cybersecurity standard bridging IEC 61508 and IEC 62443. A further concern of us is to keep this notion in line with the developments in other e.g. domain specific standards where we are active (e.g. automotive cybersecurity engineering, as explained later).

But this was not the only resulting process towards rather holistic approaches in IEC TC65 and SC65A: In the recently established new ad-hoc working groups of IEC TC65 (Industrial process measurement, control and automation) AHG2 (Reliability of Automation Devices and Systems, meeting at AIT 1.-3.6.2016, Vienna, AT) and AHG3 (Smart Manufacturing – Framework and System Architecture, kick off meeting 4.4.-6.4.2016, Frankfurt, DE, a follow-up meeting again in Frankfurt from 11.-14.10.2016). Complementary, IEC SC65E (Devices and integration in enterprise systems) started with an ad-hoc group AHG1 (Smart Manufacturing Information Models), covering the aspects of information models for exchange in context of enterprise systems, which has some impact on the work in IEC TC65A AHG3.

At these recently started ad-hoc groups EMC² and ARROWHEAD were presented as key projects to drive forward work in these new standardization topics. Since Standardization in the field of machinery (except the electro-technical aspects) is done in ISO TC 184, just now is the voting for a new work item in a joint working group ISO/IEC JWG 21 "Smart Manufacturing – Reference Models" between IEC TC 65 and ISO TC 184, which is supported by several countries of EMC² partners.

For the EMC² WP 11 somehow considering "Internet of Things" it may be of interest that the ISO/IEC JTC1 (Joint technical Committee 1) has just now decided to propose foundation of a new Subcommittee ISO/IEC JTC1 SC41 "Internet of Things and related Technologies". This efforts are undertaken to avoid duplicate standards, overlapping standards and unclear, diverting terminology – the overall (international) standardization landscape is already very fragmented and diverted, so every attempt for co-operation is supported by research and company members. These activities to align standardization with the new technology and market developments (emerging and disruptive) can be only be judged positively, and hopefully will be successful in simplifying and aligning the diverse issues arising, particularly in disruptive cross-domain areas like IoT, Cloud Computing, Smart Anything, where always many existing standards apply, but not in a holistic manner.

The "Human factors and functional safety" group IEC TC65 WG17 successfully restarted with a new convenor, Mr. Schaub, IABG, in Munich (Ottobrunn) from 4.-5.10.2016. The intention is now to write a TR (Technical Report) instead of a TS (Technical Specification) because this is easier to accomplish and finalize since so much time was lost through the "Interregnum". This report should be fed into the IEC 61508 update cycle for Ed. 3.0 (or later), so it made sense for EMC² partners who are involved in IEC 61508 Ed. 3.0 to take part.

The maintenance cycle for IEC 61508-3 (Software) started in a "preparatory mode" already two years ago because so many software paradigms arose in the meantime which are already used in safety-critical systems' development but not covered by existing standards (or even quasi "forbidden"). The Hardware- und systems' people were not so eager to start (Part 1 and 2), but are impacted by some of the proposed changes in IEC 61508-3 as well (because in many cases the system aspect is most important, not just software or hardware). Some concepts of EMC² (WP6), like contract-based development, run-time certification and guidelines or mandatory requirements to achieve security-aware safety have already been brought into the maintenance cycle as topics. The next meeting will be Nov. 30 – Dec. 2, 2016, in Vienna with AIT as host. Originally, it was planned to have a joint meeting with IEC 61508-1/2, but that was

delayed. The last meeting was in Helsinki, Sept. 6-8, 2016.

In the automotive domain, with the start of the work on Ed. 2 (ISO 26262: 2018) of the functional safety standard for road vehicles, the scope was extended to trucks, buses and motor cycles (part 12), a completely new part 11 was developed for semiconductors. This strengthens the position of EMC² addressing the challenges of new semiconductor chips and their assessment and analysis in system context. Proposals were forwarded to ISO TC22 SC32 WG08 (ISO 26262) by AIT concerning the need to address the interaction between safety and security as part to mitigate the potentially dangerous cybersecurity impact on safety. There are significant changes and additions.
Unfortunately, an important part on "SotiF – Safety of the intended Functionality" was not included in the current Ed. 2.0 (2018) version, but will be continued as a PAS (Publicly Available Specification) to be later integrated into ISO 26262 (Ed. 3.0?). This part would be most important for self-driving autonomous vehicles because a system may, without failure of a component or constituent in the sense of ISO 26262 Ed. 1.0, fail in its intended functionality e.g. by wrong interpretation of sensor input, unexpected environmental conditions and wrong perception of a situation, but was considered for a majority of national mirror committees not mature enough at the moment to become already now part of ISO 26262:2018. First results of EMC² (WP6), particularly to consider security impact on safety and to create interaction points between safety and security during the V-Life Cycle Processes (Part 2, Part 4) are already included in the current DIS. Some particular EMC² ideas, on the other hand, may not find its way in the current ISO 26262 standard.

It was, of course, not the intention to write a separate automotive cybersecurity standard in context of ISO 26262, but provide the requirements and interaction points. Details on cybersecurity counter- and mitigation measures will be handled by a complementary standard or guideline (like SAE J3061), which, on the other hand, should refer to the safety requirements and interdependencies in an adequate manner.

The idea to complement ISO 26262 with an Automotive Cybersecurity Standard was already born some time ago: so has SAE (USA) developed a guideline J3061 "Cybersecurity Guidebook for Cyber-Physical Vehicle Systems", following quite well the ISO 26262 life cycle model in all phases and distinguishing between a process for safety-related issues of cybersecurity and non-safety-related cybersecurity issues (e.g. privacy). Then have DVA (German Car Manufacturer Association) and DIN proposed a New Work Item on "Road vehicles – Automotive Security Engineering", which was focusing mainly on Security and rather separating from safety and the relation to ISO 26262 functional safety. After publication of this proposal, SAE proposed a similar NWP "Road vehicles – Vehicle Cybersecurity Engineering", based on the existing SAE Guideline J3061 (Cybersecurity Guidebook for Cyber-Physical Vehicle Systems) and referencing very clear ISO 26262. Both proposals were accepted.

This was an unusual situation, so an agreement recently signed between ISO and SAE came into force by founding a joint working group ISO/SAE JWG1 who should develop an agreed standard, now called "Road vehicles – Cybersecurity engineering" (in ISO: ISO TC22 SC32 WG11). This type of co-operation of such two totally different organized groups with different voting- and decision mechanisms is absolutely new. The kick-off meeting was in Munich, 19.-21.10.2016. Our goal was, as mentioned above, to keep developments in line with other cybersecurity/safety standards and co-engineering. First ideas collected in brainstorming groups at the kick-off meeting look quite promising for a holistic solution which would conform somehow with the results of EMC² and ARROWHEAD (and related ARTEMIS/ECSEL projects) – at least for us and at the moment. EMC² partners like AVL and AIT were active at this meeting.

As mentioned above already in context of WP 6 tools, the IOS (Interoperability Specification) is an important result of the High-Rel Cluster of projects. CP-SETIS, a support-action type IA of H2020, will organize a sound basis for maintaining and further supporting the IOS (Interoperability Specification for Tools Interoperability) and is managed by EMC² members. A second outcome will the update of the ARTEMIS Strategic Agenda for Standardization, which will be extended by (1) IOS support and (2) to include all new paradigms and developments in critical CPS and CPSoS (some of them listed above) including EMC² results.

[1] E. Schoitsch, "Autonomous Vehicles and Automated Driving: Status, Perspectives and societal Impact", IDIMT 2016, 24th Interdisciplinary Information Management Talks, Podebrady, CZ, 7.-9.9.2016; in: IDIMT Proceedings, University Series Informatik 45, Trauner Verlag 2016, Traun, AT, p. 405-425, ISBN 978-3-99033-869-8

[2] Christoph Schmittner, Zhendong Ma, Carolina Reyes, Oliver Dillinger, Peter Puschner, „Using SAE J3061 for Automotive Security Requirement Engineering", DECSoS Workshop at SAFECOMP 2016, Trondheim, NO, 20.9.2016, in: WS-Proceedings Springer LNCS 9923, p. 157-170, ISBN 978-3-319-45479-5.

[3] C. Schmittner, Z. Ma, and P. Smith. "FMVEA for Safety and Security Analysis of Intelligent and Cooperative Vehicles." 1st International Workshop on the Integration of Safety and Security Engineering at SAFECOMP'15, Firenze, Italy, September 2014

[4] C. Schmittner, Z. Ma, E. Schoitsch, T. Gruber. "A Case Study of FMVEA and CHASSIS as Safety and Security Co-Analysis Method for Automotive Cyber-physical Systems", 1st Cyber-Physical System Security Workshop (CPSS 2015) help in conjunction with ACM AsiaCCS'15, Singapore, April 2015

[5] C. Schmittner, Z. Ma, and T. Gruber. "Standardization Challenges for Safety and Security of Connected, Automated and Intelligent Vehicles." The 3rd IEEE International Conference on Connected Vehicle and Expo (ICCVE), Vienna, Austria, November 2014

[6] IEC TC65 WG20, "Framework to bridge the requirements for Safety and Security"

[7] IEC TC65 AHG2, "Reliability of Automation Devices and Systems"

[8] IEC TC65 AHG3, "Smart Manufacturing – Framework and System Architecture"

[9] IEC SC65A AHG1, "Smart Manufacturing – Smart Manufacturing Information Models"

[10] IEC TC65 WG 17, "Human Factors – Functional Safety"

[11] ISO TC 184/IEC TC65 JWG 21, "Smart Manufacturing – Reference Models"

[12] ISO TC22 SC 32 WG 8, Road Vehicles – Functional Safety

[13] ISO/SAE JWG1 Road vehicles – Cybersecurity Engineering

[14] ISO TC22 SC32 WG11, Road Vehicles – Cybersecurity Engineering

XII. EMC² Workshops and Special Sessions at renowned Conferences in 2016: CPS-Week, SAFECOMP, Euromicro DSD/SEAA

Erwin Schoitsch, AIT

In 2016, EMC² was present at and co-hosting several renowned conferences in Europe. This does not only include display of posters, folders and the like (which was done at many more conferences and workshops), but particularly presenting achievements to industry and academia.

CPS – Week in Vienna, Imperial Castle in the City Centre, April 11-14, 2016

In April an outstanding event took place in Vienna, the CPS Week. It is the leading Embedded Systems/Cyber-physical Systems Conference. It brings together four top conferences, HSCC, ICCPS, IPSN, and RTAS and numerous workshops and side events. Altogether the CPS Week program covers a multitude of complementary aspects of CPS, and reunites the leading researchers in this dynamic field.

CPS Week 2016 took place from April 11 to April 14, 2016 in the magnificent Hofburg Palace, Vienna, Austria, the former Habsburg Emperor's Palace. This multi-conference event was co-organized by Vienna University of Technology, IST Austria, AIT Austrian Institute of Technology, University of Salzburg and TTTech, with Prof. Radu Grosu from Vienna University of Technology and Prof. Thomas Henzinger from IST Austria serving as the

event chairs. Apart from the four main conferences, CPS Week hosted 21 workshops, 6 tutorials, 3 summits and the localization competition.

CPS Week 2016 was collocated with the ARTEMIS-IA Spring Event 2016, making it particularly important for the ARTEMIS community as a leading, industry driven organization in embedded intelligence and systems research. In total, it was bringing together more than 1000 researchers, students and practitioners from all around the world. The event attracted participants from 5 continents, representing 484 universities, research institutes and companies interested to discover the latest trends in cyber-physical systems research.

EMC² organized a very successful full day side-event, the so called EMC² Summit, organized by AIT on behalf of EMC2, with 17 presentations plus one external key note speaker from the University of Naples, and a full set of project posters in the adjacent hall. The papers have been collected and will be published in an open access repository as ERCIM Proceedings (European Research Consortium for Informatics and Mathematics).They provide a very good overview over the multitude of issues tackled in EMC²! The EMC² Summit ranked among the top 3 side events of CPS-Week!

The program comprised the following talks:

09:00 Welcome and Introduction
E. Schoitsch (AIT)

09:10 KEYNOTE:
Model Driven Engineering of Critical Systems
Prof. Stefano Russo, University of Naples

09:50 The EMC² Project on Embedded Microcontrollers - Progress After Two Years (Poster)
(Werner Weber, Thomas Söderqvist, Albert Cohen, Elena Garcia Valderas, Sergio Saez, Juan Carlos Pérez-Cortés, Xing Cai, Björn Nordmoen, Hans Petter Dahle, Michael Geissel, Jürgen Salecker, and Pavel Zemčík)

10:05 The S3P Project and S3P Alliance – an IoT Opportunity (Safe, Smart and Secure Platform)
(Chahinez HAMLAOUI (Systematic-Paris-Region, European Cluster of Excellence))

AUTOMOTIVE and MOBILITY SESSION

10:30 Process- and Product-based Lines of Argument for Automotive Safety Cases
(Helmut Martin, Martin Krammer, Robert Bramberger (VIRTUAL VEHICLE Research Center, Graz, Austria), Eric Armengaud (AVL List GmbH, Graz, Austria))

11:00 Implementation of active safety system for pedestrian detection in Volvos cars and real benefit of the system based on selected real-life fatal pedestrian accidents
(Peter Vertal, Gustav Kasauicky (University of Zilina, Slovakia), Hermann, Steffau (Dr.Steffan Datentechnik, Linz, Austria))

11:30 Seamless tool integration in an automotive use case: An experience report
(Andrea Leitner, Christian El Salloum, AVL List GmbH, Graz, Austria)

12:00 SOA Real-time System Development – An Automotive Case Study (Poster)
(Cuong M. Tran, Kung-Kiu Lau, Simone Di Cola (University of Manchester, UK))

12:15 Embedded Intelligence in Smart Cities through Urban Sustainable Mobility-as-a-Service: research achievements and challenges
(George Dimitrakopoulos, George Bravos and Ilianna Stampologlou, (Harokopio University of Athens))

AEROSPACE and RAIL

13:45 Implementing mixed-critical application on next generation multicore aerospace platforms (Poster)
(F. Federici, V. Muttillo, L. Pomante, G. Valente (University of L'Aquila, Italy), D. Andreetti, D. Pascucci (Thales Alenia Space, Rome, Italy))

14:00 A comparison between Hardware and Software Solutions for Resource Partitioning in Multicore-based Mixed Criticality Applications (Poster)
(Stefano Esposito, Sehriy Avramenko, Massimo Violante (Politecnico Tdi orino, Italy), Marco Sozzi, Massimo Traversone (Selex ES, Nerviano, Italy), Marco Binello, Marco Terrone (Alenia Aermacchi, Torino, Italy))

14:15 A Model-Based ESL HW/SW Co-Design Framework for Mixed-Criticality System (Poster)
(F. Federici, V. Muttillo, L. Pomante, P. Serri, G. Valente (Università Degli Studi Dell'Aquila, L'Aquila, Italy))

14:30 Deterministic Parallel Programming for Railway Applications
(Oscar Medina Duarte, Peter Tummeltshammer (Thales Austria, Vienna))

SECURITY

15:30 RoViM: Rotating Virtual Machines for Security and Fault-Tolerance
(Dorottya Papp, Levente Buttyan (Budapest University of Technology and Economics, Hungary), Zhendong Ma, AIT Austrian Institute of Technology, Austria))

16:00 Survey on Camera based Communication for Location-Aware Secure Authentication and Communication
(Hannes Plank, Thomas Ruprechter, Gerald Holweg, Norbert Druml (Infineon Technologies Austria AG, Graz, Austria), Christian Steger (Graz University of Technology, Austria))

CONCURRENCY, MULTI-PROCESSING and DYNAMIC SYSTEMS

16:30 Asymmetric Multiprocessing on industrial ZYNQ board with HDMI I/O
(Jiri Kadlec, Zdenek Pohl, Lukas Kohout (UTIA, Prague, Czech Republic))

17:00 Mining Concurrency Bugs
(Paolo Ciancarini, Francesco Poggi, Davide Rossi (University of Bologna, Italy), Alberto Sillitti* (Center for Applied Software Engineering, Italy), (*Consorzio Interuniversitario Nazionale per l'Informatica, Italy))*

17:30 Towards safe mixed critical embedded multi-core systems in dynamic and changeable environments (Poster)
(Christoph Dropmann, Tiago Amorim, Daniel Schneider (Fraunhofer IESE, Kaiserslautern Germany), Alejandra Ruiz (ICT-European Software Institute Division, TECNALIA. Derio, SPAIN))

17:45 Profile Driven Application Parallelization (Poster)
(Imran Ashraf, Nader Khammassi, Koen Bertels (TU Delft, The Netherlands))

18:00 Welfare - End of the EMC² SUMMIT

Figure 32: Impressions from the CPS Week in Vienna, Imperial Castle, April 11-14, 2016

EMC² co-hosting the DECSoS Workshop at SAFECOMP 2016 in Trondheim

Since SAFECOMP was established in 1979 by the European Workshop on Industrial Computer Systems, Technical Committee 7 on Reliability, Safety and Security (EWICS TC7), SAFECOMP has contributed to the progress of the state-of-the-art in dependable application of computers in safety-related and safety-critical systems. SAFECOMP is an annual event covering the state-of-the-art, experience and new trends in the areas of safety, security and reliability of critical computer applications. For many years, the annual SAFECOMP conferences are complemented by a series of workshops.

The DECSoS workshop (ERCIM/EWICS/ARTEMIS Workshop on "Dependable Embedded and Cyber-physical Systems and Systems-of-Systems") at SAFECOMP follows already its own tradition since 2006. In the past, it focussed on the conventional type of "embedded systems", covering all dependability aspects (in the meaning of IFIP WG 10.4, defined by Avizienis, Laprie, Kopetz, Voges and others). To emphasize more the relationship to physics, mechatronics and the notion of interaction with a somehow unpredictable environment, the terminology changed to "cyber-physical systems".

The DECSoS Workshop served during the last years as one particular dissemination event for ARTEMIS/ECSEL type projects. It is a full day workshop and co-hosted by several ARTEMIS and ECSEL projects, many of them being represented by a presentation, or posters and flyers. Already in the introduction are explained in short the co-hosting projects, to show their impact for the CPS research field in context of the European research and Innovation Initiatives in the field of CPS, among them in a prominent position EMC². The workshop was very well attended (more than 25 participants, additional chairs had to be brought into the room), and good discussions were a follow up of the interesting topics.

The Workshop had four sessions, with in total 14 presentations:

- Analysis, Test and Simulation
- Automotive
- Safety and Cyber-security Analysis and Co-Engineering
- Embedded Systems' Industrial Applications

In the introduction on European Research Initiatives in the area of CPS, in the Automotive session (two EMC² related papers) EMC² was represented, other talks covered CRYSTAL, AMASS, R5-COP, ARROWHEAD, DREAMS (FP7), several national projects and industrial case studies.
It is important to note that the main conference SAFECOMP 2016 and the Workshops have separate proceedings as Springer LNCS 9922 (main conference) and LNCS 9923 (SAFECOMP Workshops).

Euromicro 2016 DSD/SEAA in Limassol, Cyprus:

The well-established Euromicro conferences have two joint conferences which form the back bone of the Euromicro activities. These are SEAA (Software Engineering of Advanced Applications) and DSD (Digital Systems Design), SEAA was taking place for the 42nd time 2016, for DSD pi was the 19th event.

EMC² was represented twice (besides flyers and folders): In the main stream conference DSD in a session on relevant European research projects with a general presentation by our co-ordinator, Werner Weber, providing a broad overview over EMC² as a project and highlighting all major achievements of the work packages, technology as well as use cases and demonstrators.

In a special session, organized by Erwin Schoitsch, on "Software and Education Ecosystems", the ARTEMIS Approach to sustainable innovation eco-systems was presented (besides the activities like the European workshops on Education, Training and Skills, addressing the interests of the European Electronic Leaders Group). This included examples from ARTEMIS projects MBAT, SafeCer, and particularly CRYSTAL and CP-SETIS with respect to standardization and education needs to address standardization. EMC² was mentioned as one project with particular impact on standardization.

In 2017, from August 30 – Sept. 1, Euromicro will be organized by TU Vienna, OCG (Austrian Computer Society) and AIT in Vienna at TU Vienna. This will be a very good opportunity to present final results of EMC²! If there are some good ideas, we could even organize a special session for EMC². The proceedings are published by IEEE CPS and indexed by IEEE, the participants receive proceedings in an electronic version.

EMC² Impressum

Project:

- EMC² - Embedded Multi-Core systems for Mixed Criticality applications in dynamic and changeable real-time environments
- ARTEMIS Grant Agreement Number: 621429

Project Coordinator:

Infineon Technologies AG
Dr. Werner Weber
IFAG BEX RDE RDF
Am Campeon 1-12
85579 Neubiberg, Germany
Contact: werner.weber@infineon.com

Editorial:

Albert Cohen, INRIA
Luis Miguel Pinho, CISTER
Contact: emc2_newsletter@list.emdesk.eu

Layout, Design:

Alfred Hoess, Administrative Support