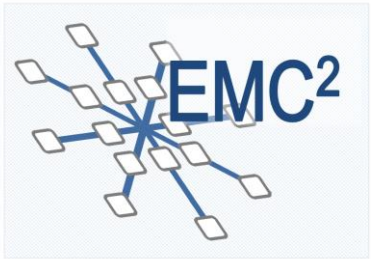



# Embedded multi-core systems for mixed criticality applications in dynamic and changeable real-time environments

**Project Acronym:**

# EMC²

## Grant agreement no: 621429

| | | |
|---|---|---|
| **Deliverable no. and title** | **D4.3 – First report on Cores, Memory, Virtualization, Accelerators, Interconnecting Cores, Chips, and Peripheral devices in Mixed-Criticality Systems: Requirements, Concepts and Preliminary Designs** | |
| **Work package** | WP4 | SP_Multi-core Hardware Architectures and Concepts |
| **Task / Use Case** | T4.2, T4.3 | Advanced processor concepts and architecture definition<br>Core and chip interconnect, peripheral devices |
| **Subtasks involved** | ST4.2.1, ST4.2.2, ST4.2.3, ST4.2.4, ST4.3.1, ST4.3.2, ST4.3.3 | |
| **Lead contractor** | Infineon Technologies AG<br>Dr. Werner Weber, mailto:werner.weber@infineon.com | |
| **Deliverable responsible** | 02E TUW<br>Haris Isakovic, haris@vmars.tuwien.ac.at | |
| **Version number** | v1.0 | |
| **Date** | 26/03/2015 | |
| **Status** | Final | |
| **Dissemination level** | | |

# Authors

| Partici-pant no. | Part. short name | Author name | Chapter(s) |
|---|---|---|---|
| 01A | IFAG | Wolfgang Schneidt<br>Werner Weber | 2.1.5, 2.6.1, 3.2.3 |
| 01R | NXPGE | Christopher Nigischer | 3.2.2 |
| 01V | TUBS | Rolf Mayer | 2.1.3 |
| 01W | TUDO | Lilian Tadros | 2.2.1, 2.5.3, 3.3.1 |
| 02A | AVL | Peter Priller | 3.2.2 |
| 02D | TTT | Andres Eckel | 2.1.4, 3.2.1 |
| 02E | TUW | Haris Isakovic | 3.1.3 |
| 04D | UTIA | Zdenek Pohl | 3.1.5 |
| 02C | IFAT | Stefan Berger<br>Norbert Druml | 2.1.2, 2.4.1, |
| 11D | Selex | Massimo Traversone | 2.5.1 |
| 15E | Tecnalia | Alonso Muñoz, Asier | 3.3.2 |
| 15O | SevenS | José Luis Gutiérrez | 3.2.1 |
| 16A | Chalmers | Ioanis Surdis | 3.1.2 |
| 17B | Imec-NL | Topbias Gemmeke | 2.2.3 |
| 17C | NXPNL | Hamed Fatemi | 2.5.2 |
| 18H | UoBr | Simon Hollis | 2.1.1 |
| 15F | TASE | Manuel Sanchez Renedo | 2.1.6 |

# Document History

| Version | Date | Author name | Reason |
|---|---|---|---|
| v0.1 | 27/04/2014 | Alfred Hoess | Initial Template |
| v0.6 | 25/02/2014 | First Draft | Haris Isakovic |
| v0.6.1 | 26/02/2014 | Internal Review | Norbert Druml |
| v0.8 | 11/03/2014 | Haris Isakovic | Final Draft |
| v0.8.1 | 23/03/2014 | Mark Honman | Internal Review |
| v0.8.2 | 23/03/2014 | Albrecht Mayer | External Review |
| v0.9 | 25/03/2014 | Haris Isakovic | Reviewed Final Draft |
| v.1.0 | 26/03/2014 | Haris Isakovic | Final |
| | 26/03/2014 | Alfred Hoess | Final editing and formatting, deliverable submission |

## Publishable Executive Summary

The deliverable D4.3 is the first technical report for the tasks T4.2 and T4.3 in WP4. It provides an overview of preliminary designs and concepts for the cores, memories, interconnects, QoS tools and methods, dynamic reconfiguration of hardware and peripheral devices.

# Table of Contents

# List of figures

# 1. Introduction

## 1.1 Objective and scope of the document

The document provides an overview of the technical work performed by the partners in WP4, specifically tasks T4.2 and T4.3. It includes a short description of a motivation behind the activity of each partner, a preliminary specification or design, planned development phases, and expected results of each individual contribution.

## 1.2 Structure of the deliverable report

The document is organized according to the structure of the work package and the corresponding tasks. First section provides introduction in the document. In the second section a report from the task T4.2 grouped by subtasks. Next section provides the report from the task T4.3 grouped by subtasks.

# 2. T4.2 Advanced processor concepts and architecture definition

## 2.1 ST4.2.1 Core-based support for mixed-criticality applications

### 2.1.1 University of Bristol (18H UoBR)

**Swallow many-core system**

Swallow is a scalable many-core system currently comprising up to 480 timing predictable processors and an on- and off-chip network.

Swallow was designed to see how the challenges of both hardware and software development develop as many-core systems scale to hundreds and thousands of processors. It has been created explicitly as a system that uses timing predictability of the processors as a mechanism to allow scalability of applications by using this predictability to reduce the need for synchronization. Swallow is also designed to expose the energy consumption of running applications, to lead to analysis and optimization.

Swallow removes the problems of shared memory accesses by using a distributed memory architecture, whereby each processor has an individual timing predictable memory. Static analysis can be performed in the input programs to determine run-time behavior and execution timing. Swallow features such as memory size, number of cores and interconnect implementation are configurable, allowing a range of systems to be synthesized.

The Swallow hardware has been completed and is on version 3. A 480 core system has been composed and can be split into a number of smaller systems. An operating system capable of supporting multiple run-time applications has been created and this is in alpha release. The OS needs extending as part of WP4 to allow the insertion of guarantees for mixed criticality applications when performing its task-to-core mapping.

The following design documents are available:

- Swallow System Overview
- Swallow nano-Operating System Overview

Swallow addresses the following WP4 requirements: TUW-001, TUW-009, TUW-010, TTT-003, TTT-010, TUDO-005, IFX-003, IFX-005, IFX-006, IFAT-002, IFAT-003, IFAT-006, AVL-001, C3E-005.

**Expected outcome**

The completion and use of Swallow will allow us to explore the problems of composition of mixed criticality task on a many-core system. We will be able to explore how tasks interact when sharing a processor and advance our understanding of timing and interaction analysis to a level that allows us to move away from needing separate processors for each critical task.

Swallow will afford a unique large-scale exploration framework to test ideas about hardware architecture, optimizations to on- and off-chip many-core interconnects and to see the interaction of test mixed-criticality workloads when running on a range of possible many-core implementations.

### 2.1.2 **Infineon Austria AG (02C IFAT)**

**Hardware/Software Multi-Core Architecture for ToF 3D Imaging**

Time-of-Flight (ToF) is a 3D imaging technology that provides distance information by measuring the travel time of the emitted and reflected light. Because of its robustness and accuracy, the ToF imaging technology is conquering the application fields of industrial, automotive, and consumer electronics. However, the 3D algorithms involved require considerable amounts of computation power. Therefore, novel hardware/software approaches are developed in order to cope with todays and future challenges.

The design approach focusing on the automotive domain uses a state-of-the-art automotive multi-core controller. While 3D processing shall run on one core, 3D data interpretation and other tasks shall run independently on other cores. In addition, tailored and hardware-accelerated co-processing units shall be used in order to speed up computation-intense calculations.

During the first year, requirements are developed based on the identified use-case scenarios. After these requirements and specifications ware defined, a concept for a multi-core hardware/software architecture is developed based on extensive research activities. During the second year, the proposed hardware/software design will be implemented. Finally, demonstrators will be realized and will showcase the advances achieved in ToF processing. Furthermore, detailed evaluations in terms of computation performance, memory requirements, and power consumption will be carried out.

The projected demonstrator is described in more detail in deliverable D4.15. The detailed requirements and concept report is currently in work and will be available through deliverable D4.21.

**Expected outcome**

The architecture and technology developed during this project will enable ToF 3D processing in an efficient and parallelized way, which will be exploited for anticipated applications in the field of automotive and consumer electronics.

The architecture and demonstrators developed during this project will feature state-of-the-art techniques in order to enable multi-core ToF 3D processing in a mixed-critical environment, which is especially important for the automotive domain.

The concepts and technologies developed during the course of this project represent novel solutions for ToF 3D imaging, in particular in the fields of automotive and consumer electronics. Furthermore, the know-how and knowledge base will grow and will reach out to promising fields for future market applications.

### 2.1.3 **TU Braunschweig (01V TUBS)**

**Enhanced Run-Time Monitoring for Mixed-Criticality**

On a multi- or many-core platform that runs applications of different safety criticality (mixed-criticality), all applications have to be certified to the highest level of criticality, unless they are sufficiently isolated. Isolation enables individual certification of applications and cost-efficient re-certification of single applications after an update. We introduced a parameterizable and synthesizable many-core platform with a fast and scalable monitoring and control mechanism that supports safe sharing of resources.

With the IDAMC we introduced a parameterizable and synthesizable many-core platform that allows the evaluation of concepts and mechanisms for running multiple mixed-critical applications on a single platform. But for further development, we need the design on a higher abstraction to enable faster

evaluation of new parameters and monitoring and control capabilities. So we re-implement the design in SystemC based on the SoCRocket platform and extend it with an interface for easy addition of such.

**Development phases**

- Step 1: M12 (end of March 2015)
  - A concept defined
  - Initial SystemC NoC Model
  - The base architecture defined
- Step 2: M24 (end of March 2016)
  - Tested NoC Model
  - Initial network interface with parameters and capabilities
  - Initial prototypical integration of EMC2 results demonstrated
- Step 3: M36 (end of March 2017)
  - Detailed performance evaluation exists
  - Final EMC2 demonstrator available and accessible to whole consortium
  - Demonstrates innovations and input from EMC2

Contributing to high-level requirements as defined in Requirements_Document_D0401. Overview defined in D4.2 Specification of Demonstrator Platforms (chapter "Communication and verification of mixed-criticality MC control units"). Additional resources can be found in [1], [2], [3].

**Expected outcome**

We show the effects of sharing resources by applications of different criticality on the response time of a highly critical task. With our monitoring and control mechanism, we develop a fast and scalable solution to isolate the timing of mixed-critical applications to reduce their cost for certification and re-certification. We demonstrate the effectiveness of the mechanism to isolate a highly critical task against a (faulty) low criticality task.

### 2.1.4 TTTech Computertechnik AG (02D TTT)

**Portability of the ACROSS MPSoC architecture to a Xilinx Zynq architecture**

This activity focusses on the TTT cooperation with TUW in enhancing the results jointly developed in the frame of the ARTEMIS JU project ACROSS. In the frame of this project a Multi Processor System on a Chip (MPSoC) was developed using the GENESYS architectural template for embedded systems design. The results were based on the ALTERA Stratix IV FPGA and eight NIOS cores. The aim of the project was to demonstrate the multi industrial domain usability of the architectural approach and the advantages for certification and re-use compared to classical approaches. The project did not aim at maximizing performance compared to the available technical options with respect to high performance cores and devices. Introducing a set of high performance cores would have exceeded the ACROSS budget limits and also would have caused significant problems in terms of usable space in the FPGA selected. Thus, prior to significantly investing in a high performance solution, the NIOS cores were accepted as a feasible compromise in order to investigate aerospace certifyability and the advantages of using the GENESYS architectural template in multiple industrial domains. The results in INDEXYS and ACROSS demonstrators were strongly encouraging to carry on activities and to enhance the development towards high performance solution even more since in the meantime two suitable platforms were launched on the market: the Xilinx Zynq and the ALTERA ARRIA platforms offering ARM-cores plus FPGA space for integrating new IP on one device.

Thus the work first of all will focus on cooperating with TUW in terms of porting the results to the ALTERA ARRIA Platform in order to gain the desired performance increase compared to the NIOS cores

used in ACROSS. This shall show if the performance of the MPSoC is mainly dependent on the cores used or if other factors are limiting or influencing the performance.

Furthermore, it is envisaged to also look at the Xilinx Zynq platform and investigate if this platform would have benefits compared to the ALTERA ARRIA approach.

**Expected outcome**

We expect that the results gained during the ACROSS project will significantly be enhanced by using (a) the ALTERA ARRIA platform (b) the Xilinx Zynq platform compared to the performance possible using the approach employing the ALTERA Stratix IV and the NIOS cores. We further are confident that the architecture as such is independent of the platform and so does not limit the performance. We expect that the performance of the approach is mainly dependent on the performance of the cores used. In addition, we therefore target to demonstrate that the enhanced development will allow the approach to be used for more sophisticated applications and thus will pave the way to introduce the results of GENESYS, INDEXYS, ACROSS and EMC2 in highly safety relevant applications in many industrial domains. We already know from the ACROSS project that the approach is certifiable according to aerospace standards up to DAL A. This will provide an option to the aerospace industry to either use the approach in an FPGA implementation or to introduce the NOC in modified or even specific (ASIC) implementations in order to solve the inherent problem of using microprocessors for future aerospace applications since certification causes problems when the internal NOC is not open for review by the authorities. We see that similar problems will arise in other domains such as the automotive domain and when looking at highly automated driving and potential certification demands and complexity issues coming up as well. Today it is crucial that the NOCs of COTS multicore components cannot be investigated by the certification authorities due to IP reasons (manufacturers do not open NOC IP for review to authorities).

Thus we expect to contribute significantly by these activities to solve the problem of multiple industries with respect to their future products requiring certification in highly safety relevant applications.

Demonstration is envisaged in WP 7 (automotive) and WP 9 (space).

### 2.1.5 Infineon Technologies AG (01A IFAG)

**Tracing Mixed-Criticality**

Integrating software with mixed criticality, specifically different ASIL levels onto one microcontroller is feasible by avoiding interference in time and space. Modern microcontrollers provide different means for avoiding interference, going beyond a simple MPU and timers. However, all the means need to be configured properly. The objective of this sub working package is to provide both a methodology and a tool, checking based on a non-intrusive on chip trace system, whether the microcontroller was configured properly to avoid interference in space and time. Some examples for interferences are listed below:

- Incorrect MPU configuration, allowing for joint write data access between different criticalities
- Interference on joint bus systems, for example a peripheral bus, between different cores with different criticalities

The tool will not only use the on chip trace system, but also check information which is contained in the elf file, the linker script and possibly even parts of the source code.

**Development phases**

- Step 1: M12 (end of March 2015)
    - Tool and methodology concept defined

- Step 2: M24 (end of March 2016)
  - Tool prototype implementation as Eclipse RCP
- Step 3: M36 (end of March 2017)
  - Refined tool implementation, including different options for report generation
  - Methodology and tool assessment, introduction of the methodology and tooling on a safety related conference

**Expected outcome**

Assuring a certain safety level on a modern multi-core microcontroller is a complex undertaking. On the one hand microcontrollers in the automotive domain are under a certain cost pressure, meaning computational resources and memory are used without lots of head room. On the other hand, in particular for chassis systems with more and more advanced driver assistance systems, it is more and more needed to comply with ASIL B, C and even D. The suggested methodology and tool shall help to reach higher integrity levels in a systematic manner. Long term a qualification of the tool will be needed; however, this is beyond the scope of this research project.

### 2.1.6 **Thales Alenia Space Spain (15F TASE)**

**MPSoC hardware architecture for on-board signal processing**

The increasing demand of data proccesing algorithms in new payloads requires to design new digital hardware architectures. TASE has developed a hardware platform to integrate massive data processing MPSoC architectures in order to validate new signal processing algorithms for space applications. It has been designed to be used as a generic platform for the following space signal processing techniques: Regenerative transponders, adaptive air interface, interference cancellation/mitigation technologies, high efficiency modulation/coding techniques and spot beam technology. An application based on sub-band signal processing will be implemented to demonstrate the ability to carry out a multiband analysis.

The hardware platform is based on two space equivalent FPGAs. The first one is a Rad-tolerant ProAsic 3E from ACTEL. This FPGA will integrate a LEON2 processor to control the signal processing algorithms based on real-time operating systems like RTEMS (space-qualified OS). Signal processing algorithms will be integrated in a Virtex5 FPGA specifically useful for high-performance DSP applications. A DSP controller based on an embedded microprocessor will be used to interface with LEON2 processor and to command and control DSP applications. LEON2 and DSP controller will interact with a dual-port RAM as a mail-box. This technique is widely used as a message passing mechanics. The dual-port RAM will be integrated in the LEON2 memory map by using the specific IO signal in order to use the dual-port RAM as a specific page of the memory.

In addition, the proposed hardware platform will be used to integrate self-healing architectures proposed by Tecnalia (15E) in WP4 & WP9.

**Expected outcome**

The MPSoC architecture for massive data processing brings the advantage of a rapid hardware evaluation for data processing algorithms in space equivalent FPGAs. The collaboration with Tecnalia will allow to demonstrate the feasibility of on-board reconfiguration with LEON processor controller by using a Dual port RAM as a mail-box between processors.

## 2.2 ST4.2.2 Multi-core local and distributed memory hierarchy, interfaces, modules and compression techniques

### 2.2.1 TU Dortmund (01W TUDO)

**Deterministic and Efficient Memory Hierarchies for Mixed-criticality Real-Time Systems**

Thread execution on tightly-coupled multi-core systems with shared memory, local caches and central task scheduling shows a largely nondeterministic timing behavior due to complex cache interactions. In order to provide timing predictability, cache interactions must be restricted to related cores and must not harm cores used for other applications. Freedom from interference is a first class requirement of mixed-criticality systems.

A SystemC design-space-exploration platform is being expanded and analyzed regarding predictability for different core arrangements, memory strategies and coherence mechanisms. In detail:

- Different cores will be integrated for flexibility and for understanding the effects of heterogeneity on data integrity and analyzability.
- Cache coherence techniques will be analyzed and adapted for the best combination of throughput vs. determinism.

**Development phases**

- New core definitions and integration
  - Status: Platform currently supports multiple LEON3 cores
  - Ongoing: ARM Cortex A9 currently in integration phase
- Definition and implementation of core/memory interfaces
  - Status: LEON3 L1 cache (write-through, non-write-allocate, bus snooping)
- Definition and implementation of cache coherence mechanisms
  - Status: Implemented several cache algorithms (write-back, write-allocate)
  - Ongoing: Integration of deterministic cache coherence algorithm
- Simulation of relevant multi-threaded benchmarks
  - Status: Platform runs GRLIB tests and some standard signal processing benchmarks
  - Upcoming: Running ParMiBench multiprocessor benchmark
- Runtime analysis of coherence algorithms

**Expected outcome**

We are working towards a high-performance, dependable and analyzable core/memory architecture and algorithms with the simulation platform as proof-of-concept. This combination of reliability, determinism and performance will outstrip the state-of-the-art and is a core target of EMC² and a key ingredient in future mixed-criticality platforms.

### 2.2.2 NXP Germany (01R NXPGE)

**Memory compression techniques**

For EMC memory is a critical cost adder that needs to be minimized to achieve competitive solutions. Memory compression by means of audio source compression was investigated. As an example a digital broadcast receiver was used:

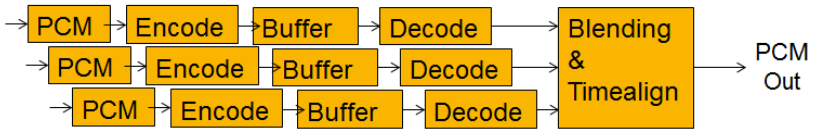


**Figure 1: Memory compression**

**Development phases**

NXP developed a technique to reduce the memory footprint of a hybrid broadcast embedded multicore receiver dramatically. Depending on the radio standard the audio time alignment buffer is the main memory consumer in a broadcast receiver. Different approaches were investigated to introduce a compression technique with the intention to reduce on chip memory. Similar to commonly known audio compression (e.g. mp3) there is a choice between memory consumption and quality / performance. A challenge is that the encoding techniques are block oriented approaches which can't be used for sample true time alignment. An intelligent buffer technique was used to allow sample based compression on top of a block based coding algorithm.

**Expected outcome**

By introducing the compressing of the audio PCM data, we achieved a memory reduction by a factor of 5. Due to the intelligent buffer techniques next to the block based audio encoder even a sample true alignment of several audio streaming sources was achieved.

### 2.2.3 Imec-nl (17B imec-nl)

**Modeling of Reliability of Memories**

With the advances in technology development and the push to reduce power consumption, memories are becoming the weak link in terms of reliability. Exploiting trade-offs between power efficiency and reliability can be achieved in a first order by using large, high-voltage memories for the critical parts and reduced footprint instances running a low supply voltage for parts of the system that tolerate errors. In order to overcome this strict separation, the reliability of memories and their failure behavior has to be modeled. Due to the complex dependency of various failing mechanisms in a memory, it is compulsory to calibrate such models with data from silicon measurement.

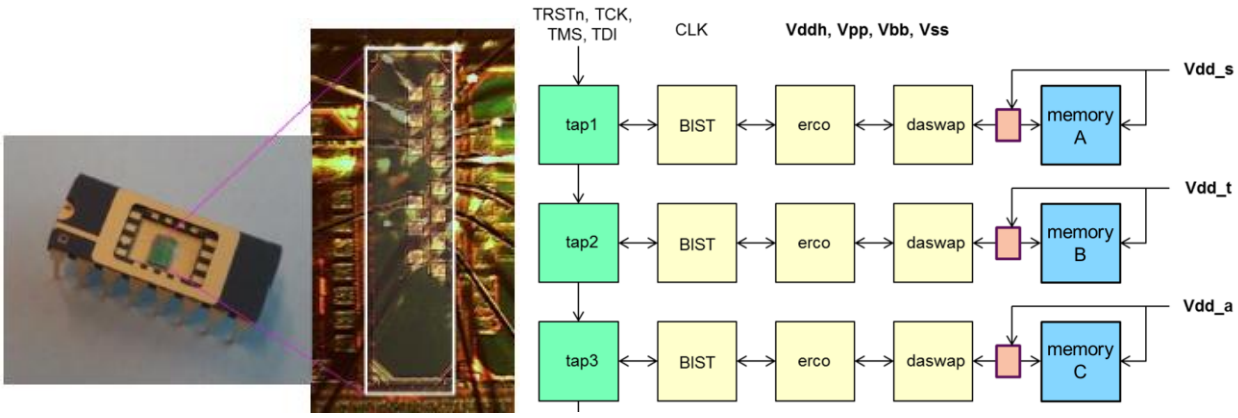


**Figure 2: Modeling memory reliability**

**Development phases**

Imec-nl has started by producing silicon samples that can be used to characterize failure behavior and energy efficiency as function of applied voltages and required performance. Based on these results we will model the behavior so it can be used for system optimization. Finally, proposals will be made for reliable embedded systems that are capable of trading-off reliability with energy efficiency while meeting hard real-time constraints.

**Expected outcome**

A memory error model calibrated with silicon measurements, which will enable the development of new error mitigation techniques that target improved power efficiency under hard real-time constraint for embedded systems. Dynamically varying systems knobs will allow to adopt reliability with energy efficiency in the system.

## 2.3 ST4.2.3 Fixed and reconfigurable hardware accelerators and co-processors

### 2.3.1 Chalmers (16A Chalmers)

**Reconfigurable RISC architecture for tolerating permanent faults**

Recent semiconductor-technology trends have resulted in the emergence of new reliability challenges in MPSoC design. Smaller transistors are more sensitive to manufacturing defects, wear-out phenomena and dielectric breakdown, all of which result in permanent faults. Even a small portion of the faulty circuitry could render an entire chip unusable, lower production yields or cause failures during the chip lifetime.

A processor of a high availability system needs to be protected and tolerate permanent faults in order to support critical tasks. In this work, we modify a RISC architecture to tolerate permanently damaged components. We divide the processor in multiple substitutable blocks and use reconfigurable interconnects to isolate them when faulty, and share them to neighboring identical processors if spare. In addition, we use blocks of fine-grain reconfigurable logic to replace faulty processor parts if a spare identical component is not available. The above require architectural changes in order for the processor to tolerate the increase delay introduced by the reconfigurable interconnects and the spare components.

In the first year, we started from an existing RISC processor, worked on the definition of the above concepts and explored their potential efficiency in terms of performance and fault tolerance. Subsequently, in the second year we will proceed with the hardware implementation of the design. Finally, in the third period of the project we will evaluate the performance, area, power and fault tolerance of the implemented reconfigurable RISC processor.

The above technique will increase the tolerance of an MPSoC to permanent faults improving its availability, while at the same rescuing damaged processors, thereby providing a more gracefully degradable system, allowing critical tasks to run on them.

**Expected outcome**

The above technique is expected to provide high tolerance to tens of permanent faults on an array of 4-8 RISC processors increasing their availability in the system. Our goal is to tolerate 2 times more permanent faults compared to a sparing technique at the processor level occupying the same silicon area.

Mixed-criticality multicore systems require better processor availability to ensure that critical tasks have functioning processors to host them. Designing such processors to be tolerant to permanent faults increases system's availability and robustness.

So far many techniques have been proposed for the processors tolerant to permanent faults (e.g. StageNet). These techniques, however, suffer from high performance overheads or low fault tolerance. In our approach we intend to pipeline the reconfigurable interconnects, hiding the wire delay. This can be achieved if the microarchitecture is able to tolerate variable pipeline depths of a processor.

## 2.4 ST4.2.4 Multi-core power management components

### 2.4.1 Infineon Austria AG (02C IFAT) and Infineon Technologies (01A IFAG)

**AMSPS-Systems for MCMC integrated Systems**

The rapidly increasing computing power in many applications and the further trend of integration of many system functions leads to new and demanding challenges for micro-controller development of various application fields (e.g. for automotive or industrial subsystems). These subsystems (e.g. motor management of vehicles) combine sensor and actuator functions with high compute power and programmability.

Increasing compute power goes hand in hand with an increasing requirement set for accompanying power supply, measurement, and support functions for the future multi-core architectures. Furthermore, the complexity of new integrated microcontroller systems with embedded multiple cores require new concepts of integrated power supply management architectures, new demanding concepts for measurement functions (e.g. ADCs,…), and other support functions (e.g. PLL, high speed interfaces,…).

Technical requirements of power supply architectures of integrated multi-core systems shall offer a variable input voltage range of different applications. Furthermore, multiple input supply concepts shall be possible to offer high application flexibility.

A Power Management System must support any possible start-up condition as long as maximum ratings in terms of input voltage and overload condition at the input pins over the ambient temperature range are not violated. Power supply levels and voltage monitoring accuracy shall fulfill the requirements for steady state accuracy to ensure the safe function of multi-core systems. Additionally, dynamic functionality for safe function of multi-core systems should be offered accordingly.

Voltage monitoring accuracy is offered to ensure a safety compliant function of multi core systems. The Power Management scheme allows activation of power down modes so that the system operates with the minimum required power for the corresponding application state.

High efficiency is required within the power management concepts for various MCMC system modes. Maximum output current of voltage regulator should support the need of multi core systems to ensure full functionality without performance reduction within the systems.

Investigations on functions to reduce the static leakage current and potentially new derived concepts are needed to keep the system performance high. Investigation on new solutions for reduced quiescent current of a multi - core system is needed (e.g. in Standby Mode).

Chip interconnect investigations according to AMSPS power supply activities are covered in ST 4.3.2 additionally. The focus lies on power supply and ground interconnect and also reference interconnect between different AMSPS functionalities (e.g. reference distribution to ADCs etc.). Furthermore, digital interconnects are investigated with the focus on functional control and data transmission to the AMSPS functionalities. This refers also to the activities of ST 4.3.2.

**Development Phases**

1) Definition and refinement of functional requirements (1-12)
2) Definition and refinement of Safety requirements (6-18)
3) Definition of AMSPS Concepts for integrated MCMC systems (1-24)
4) Definition of test chip Concepts (1-18)
5) Implementation of test chips for AMSPS functionality (1-24)
6) Verification and Validation of test chips (6-36)
7) Final documentation (24-36)

**Expected outcome**

AMSPS Concepts for integrated MCMC systems defined and partly verified on test chips.AMSPS Concepts for integrated MCMC systems with safety requirements implemented to secure ISO26262 compliance.
AMSPS Concepts for integrated MCMC systems with safety requirement implemented to secure ISO26262 compliance defined and pre-verification of new challenges of the next generation multi core systems.

## 2.5 ST4.2.5 Virtualization of cores and resources for the Virtual Hard Real-time Machine

### 2.5.1 SELEX (11D SELEX)

**Mixed Critical Avionic Application**

This activity focusses on the Selex-ES cooperation with Alenia and Polito to evaluate and test different applications with different criticality levels and different iteration rates with respect to multi core processor.

The target of this research activity is to implement avionic applications in a multicore processor to assess the benefits of such HW architecture in contrast to the classical hardware architecture used in the current avionic system. Moreover the scope of this activity is also to implement applications with different safety levels and provide mechanisms to manage such a different safety level in a real time environment and in the new HW architecture

The two applications will have different DAL because the functions provided by them are very different from the safety point of view. The most critical application will provide the data sensors monitoring and through some algorithm and filtering it provides the output data for other systems and will provide also the properly caution signal to the pilot whenever the situation (in term of elaborated data) is not stable or exceed the warning levels. The second application will provide the data conversion in term of engineering values and data for the proper visualization.

These applications will also run at different iteration rates and this makes even more challenging the design and the implementation of these applications with these new hardware and software technologies in avionic field.
The processor boards suitable for critical avionic application are always characterized by the presence of an FPGA. This item provides some “special HW functions” not built into a standard CPU. The VHDL developed for this FPGA have to be compliant to the DO254 guidelines with respect to the DAL requested in the project. The FPGA is useful for developing uncommon peripherals for input/output discrete data, for redundancy management and for security aspect (example avionic watchdog). The

presence of one FPGA is an aspect that cannot be left out for the compliances with Aircraft management applications.

The FPGA therefore will be used to evaluate the two different avionics functions and in addition a new design of the typical "Avionic watchdog function" will be implemented. This function is normally developed outside the microprocessor to assure that the SW cannot disable it. A function that can be disabled only by external events (HW) directly connected to the FPGA.

The implementation of this function is well known for a single core, but not for multicore processor with mixed critical applications. The simplest watchdog extension to the multi-core is to provide separate single watchdogs one for each core. The question is: it is correct to reset the equipment when a non-critical application fails and therefore to lose also the critical application? Probably not, and we would like to evaluate and to develop a complex watchdog suitable for the mixed-critical applications on multicore where the decision about the equipment reset is correlated to the safety critical level of the application faults. This functionality will be composed of a HW implementation and also a SW layer to manage it.

The avionic application will be implemented in a multicore processor board with FPGA embedded based on a Zynq processor (Zynq-7000 SoC). This processor contains two ARM A9 cores and also an FPGA in the same chip. It covers both the "basic" HW requirements:

- Multicore;
- FPGA presence.

**Development phases**

- Step 1: M12 (end of March 2015)
  - Concept defined
  - Base architecture defined
- Step 2: M24 (end of March 2016)
  - detailed performance
  - initial prototypical integration
- Step 3: M36 (end of March 2017)
  - final demonstrator available and accessible to whole consortium
  - demonstrates innovations and input from EMC2

**Expected outcome**

A MCP component remains however a COTS component, i.e. it features proprietary data from the COTS manufacturer: the complexity of using MCP in Mixed Critical / Hybrid avionics application, has to take into account the level of design assurance that could increment in complexity the application (including redundant or dissimilar architecture). A reduction of the complexity and difficulties that arise from the use of MCPs while meeting required deterministic behavior and target levels of performance integrity has to be analyzed/detailed closely: MCP features linked to Shared Resource Accesses like Memory, Bus, Network, Internal Registers, Clock Management, etc., must be closely analyzed.

The system design with the new HW and SW architectures will therefore be evaluated in term of performance to provide a benchmark with the same implementation on a traditional avionic hardware.

An expected outcome should be considered as "suggestion" to be discussed with the MCP manufacturer or with the certification authorities in order to understand the best approach to use MCP in a critical environment (DAL A-C).

### 2.5.2 **NXP Semiconductor (17C NXPNL)**

**Parallelization in multi-core architecture**

In the last decade or so, the design focus has shifted towards maximizing computing throughput and reducing energy consumption for new applications in the automotive domain. In this field, a central processor/microcontroller collects information from different sources (sensors or other microcontroller platforms) to make a decision/measurement. Most of this information reaches the processor as interrupts. Various techniques at architecture and circuit levels have been explored to maximize the throughput and to minimize latency of a computing platform, like increasing the frequency of computation or adding parallelization to the system. These techniques lead to an increase in the total power dissipation of the system. To compensate the increasing energy, several circuit techniques have been introduced to reduce the energy of the system like power switches, body biasing, clock gating etc.

In this project, the emphasis is on energy efficient computation while maintaining the overall system throughput. This can be achieved by minimizing the system's operating voltage in a parallel architecture geared up to balance the workload among the computational units.

**Development phases**

In this project, we will explore parallelization (techniques in a parallel system environment with intelligent scheduler. The resulting parallelization along with voltage scaling is used to exploit the overall reduction in energy consumption of the system for a given workload. This project incorporates an intelligent energy manager to balance workload in the parallelized system. The proposed research topics are:

- Modeling
    - Analysis of workload balancing in parallel architectures
- Development
    - Parallel architecture with intelligent scheduler
    - Parallel architecture with workload based Adaptive Voltage Frequency Scaling (AVFS)

**Expected Outcome**

In this project, we are making a demonstrator to show the reduction of power consumption in parallel system by using the intelligent scheduler while throughput remains constant.

### 2.5.3 **TU Dortmund (01W TUDO)**

We aim to provide a virtualization layer for abstracting and arbitrating resource accesses as described in detail in 3.3.1. The virtualization layer will transparently perform address translation of core-specific memory maps as well as regulating concurrent accesses to system resources using deterministic and criticality-aware arbitration. Apart from addressing resource sharing bottlenecks, the provided abstraction will ease the integration of heterogeneous, criticality-aware many-cores.

## 2.6 ST4.2.6 Design and Implementation of System Control for multi-core microcontrollers

### 2.6.1 Infineon Technologies (01A IFAG)

The activities in this subtask performed by IFAG support their work in other subtasks. IFAG's contribution to ST 4.2.6 concerning the inter-chip communication aspects is already reported in section 0 and the power management topics in 2.4.1.

### 2.6.2 TU Braunschweig (01V TUBS)

Our contribution to Task 4.2.6 is tightly coupled and integrated with the research and implementation in Task 4.2.1. The system control capabilities are implemented and managed via the monitoring and control mechanisms form Task 4.2.1. Due to that tight coupling we intend to follow the same milestones as in Task 4.2.1.
Moreover, the outcome will show that the same mechanisms not only help to supervise individual processing tiles, but will help to manage a network on chip in a centralized or distributed manner. Selected details of the resulting solutions will be integrated into the demonstrator platforms produced in Task 4.6.

# 3. T4.3 Core and chip interconnect, peripheral devices

## 3.1 ST4.3.1 On-chip interconnect for mixed criticality systems

### 3.1.1 Chalmers (16A Chalmers)

**Resilient Service-oriented NoC**

Technology scaling comes with great challenges. On-chip resources become more susceptible to manufacturing defects and wearout phenomena, which may damage permanently a circuit and render an entire chip unusable. In a multicore system, the Network on-chip (NoC) can be a single point of failure as it shares its resources with every component on a chip.

As a consequence, a service-oriented NoC (that is a NoC that provide particular QoS guarantees for different traffic-flows/services) needs to integrate some techniques for permanent fault tolerance. In our approach, we characterize the network resources based on the particular service they support and when faulty bypass them allowing the respective traffic class to be redirected. We work on two alternatives for service redirection, each having different advantages and disadvantages. The first one, Service Detour, uses longer alternative paths through the resources of the same service to bypass faulty network ports, keeping traffic classes isolated. The second approach, Service Merge, uses resources of other services providing shorter paths, but allowing traffic classes to interfere with each other. The remaining network resources that are common for all services employ additional mechanisms for tolerating faults, namely wire shifting and sparing for NoC links and TMR for the control-path of the Routers.

In the first year, we started from an existing service-oriented NoC and worked on the definition of the above concepts and explored their potential efficiency in terms of performance and fault tolerance. Subsequently, in the second year we will proceed with the hardware implementation of the design. Finally, in the third period of the project we will evaluate the performance, area, power and fault tolerance of the implemented resilient service-oriented NoC.

The above technique will increase the tolerance of an MPSoC to permanent faults improving its availability, while at the same preserving to the degree possible the required QoS for the system.

**Expected outcome**

The above technique is expected to provide high tolerance to tens of permanent faults on a NoC (approximately 1-100) while preserving the guaranteed packet latency and throughput required by each traffic-flow. Our goal is to maintain 50% of the network throughput, offer 2x of packet latency for high priority packets and preserve 90% of network connectivity when tolerating tens of faults.

Mixed-criticality multicore systems require NoC support for different types of traffic, thus a service-oriented NoC. Designing such NoC to be tolerant to permanent faults increases system's availability and robustness.

Currently many techniques have been proposed for fault tolerant NoCs, however, very little has been done to protect service-oriented NoCs from permanent faults and even less in describing a holistic approach that covers all NoC parts.

### 3.1.2 TU Wien (02E TUW)

TU Wien is working on an upgrade of the MPSoC based on time-triggered network on chip (TTNoC) developed in the project ACROSS.

**High-performance TTNoC Many-Core Architecture**

The evaluation of the ACROSS MPSoC identified performance as one of the essential improvement points of the architecture. The processing units used in the project were based on FPGA soft-coded solutions. As these processors worked on lower clock-rates their performance capabilities were equivalently limited. Lack of computing power prevented full evaluation of the platform. In the EMC2 project the platform will be extended by integrating a component built on a hard processor system (HPS). The performance capabilities of the given component are significantly better compared to the components used in the ACROSS MPSoC. Although, this solution is still far from a dedicated chip implementation, it will provide us with valuable insights towards this goal. The focus of the evaluation will be heterogeneity, mixed-criticality, overall performance, the performance of the TTNoC, and an integration of multi-core components.

The platform is based on a hybrid SoC architecture, featuring an FPGA and an HPS on a single chip. The FPGA will host a version of the TTNoC and a number of soft-coded components. The HPS will be connected to the TTTNoC as an independent component and it will be based on ARM Cortex A9 dual-core processor. For the implementation of the prototype the Altera Arria ST SoC development board will be used.

The work is currently being done on the preliminary design and specification, and porting of the legacy hardware and software. In the year two of the project, it is expected for a first prototype to be finished, it will be used as a proof of concept. Further, in the years two and three, it is planned to realize an industrial demonstrator based on this platform. On these tasks we are cooperating with other partners from WP4 and WP9.

The additional references and requirements for the work are defined in deliverables: D1.1, D4.1, D4.15, D4.2, D9.1. A document is currently in development called “Specification of High-performance TTNoC Many-Core Architecture”, it is a detailed preliminary design of the architecture.

**Expected outcome**

As the final result of the project we expect to have a working prototype capable of implementing industrial demonstrators. We plan to evaluate TTNoC using increased performance capabilities of the components. Also, the TTNoC MPSoC is a heterogeneous architecture and we plan to use this property to explore the idea of multi-core components. The results shall be demonstrated in WP7.

### 3.1.3 The Institute of Information Theory and Automation (04D UTIA)

**Asymmetric Multiprocessing on ZYNQ**

The goal is to provide an execution platform based on the latest generation of Xilinx devices (Zynq) with hardware acceleration. We also aim to support development on this platform by providing reference designs, examples and latest tool flow support for EMC2 partners.

The detailed platform architecture is described in D4.15 First demonstrator platforms. The integration of the design with the latest tool flow requires analysis and redesign of the on-chip interconnect network in order to integrate accelerators with the new bus standards.

**Development phases**

- Analysis of the possibilities to support mixed criticality applications in the current implementation of the UTIA Asymmetric Multi Processing (AMP) platform consisting of ARM and MicroBlaze CPUs on ZYNQ (28nm).
- Analysis of differences of auto-generated interfaces for system buses on ZYNQ (ARM AMBA3 and ASIC IP High performance ports) and the FPGA AXI4 hierarchical interconnect on the Xilinx Artix FPGA board AC701.
- Analysis how to use the Vivado High Level Synthesis (HLS) IP interfaces to AXI4 and AXI4-Lite for the integration of IP blocks to the system interconnect.
- Redesign and integration of EdkDSP accelerators in latest Xilinx tools. Update of implementation guidelines.

**Links to the requirements**

HL-REQ-WP07-005 Harmonization of development tools, HL-REQ-WP07-041 Low latencies, HL-REQ-WP07-043 Early platform available for EMC

**Expected outcome**

From the beginning of the project available execution platform which initially uses older tool flow and will be gradually improved in terms of performance and architecture. The latest evaluation version will be downloadable free for all EMC2 partners from [4]. The package also includes usage guidelines.

Novel technology in comparison to state-of-the-art:

- EdkDSP HW acceleration in AMP design.
- Critical and non-critical section support.

## 3.2 ST4.3.2 Inter chip- and system-interconnect

### 3.2.1 TTTech Computertechnik AG (02D TTT)

**TTEthernet WireLess (WL)**

The TTEthernet product capabilities shall in the long run be connectable to wireless communication systems in order to allow direct data exchange with other vehicles (V2V communication) or fixed stations such as base stations of OEMs or traffic managing organizations. In principle the current AS6802 SAE Standard TTEthernet is not limited to wired data communication. This work shall investigate and implement a possible solution to connect the TTEthernet based equipment ((a) “switches” and (b) end-systems) to wireless devices in order to support the data exchange also with other stations like other vehicles or base stations etc. Adequate development involving hardware solution, middleware support and configuration tools shall be conducted.

**Expected outcome**

The outcome of this activity will enhance the TTEthernet based IP for TTEthernet Switches and TTEthernet Endsystems by a suitable wireless interface that supports safety and security issues. The results will also provide the needed software and drivers to operate such interface appropriately and according to the TTEthernet communication principles and specifications. Such functionality will reply to the request by numerous potential customers in different industrial domains such as automotive, off-highway applications, off-shore wind power plants, aerospace and industrial robotics will close a gap in support for next generation of equipment applications. It will allow to open up the range of application opportunities of TTEthernet based equipment to connect to any other kind of station in the cloud. Thus, it

will be possible to for instance retrieve data from other vehicles or from fixed base stations. By doing so, the applications such as uploading new versions of software by an OEM on request of a client can for instance be conducted without consulting a garage. A client could order such upgrade by smart phone or similar and request uploading. Payment functions of the smart phone can be used and the OEM can then conduct the upload process, configuration and test in a remote manner.

In addition, most recent versions of information for the driver, available in the cloud, can be retrieved from the internet. The driver (or later potentially on the automated driving computer) can take them into consideration and react accordingly.

### 3.2.2 **AVL (02A AVL)**

**Networking for V&V of mixed-criticality applications on multicore automotive control units**

Multiple tasks running on multicore xCU's shall be enabled to communicate to V&V systems via unified high-performance, low latency real-time networks (CAN, CAN-FD, automotive Ethernet etc.)

Time-determinism (e.g. cyclic real-time requirements), low latency, data consistency and synchrony will be addressed in conjunction with significantly increased data throughput. Analysis and extension of automotive protocols like XCP and frameworks like AUTOSAR for above described demands providing support for high-bandwidth, multi-application, mixed criticality data streams shall be considered.

**Development phases**

- Step 1: M12 (end of March 2015)
  - Concept defined
  - detailed requirements defined
  - base architecture defined
- Step 2: M24 (end of March 2016)
  - high performance test application demonstrate-able on xCU and test system HW as specified
  - detailed performance evaluation exists
  - initial prototypical integration of EMC2 results demonstrated
- Step 3: M36 (end of March 2017)
  - final EMC2 demonstrator available and accessible to whole consortium
  - demonstrates innovations and input from EMC2

Contributing to high-level requirements as defined in Requirements_Document_D0401. Overview defined in D4.2 Specification of Demonstrator Platforms (chapter "Communication and verification of mixed-criticality MC control units"). Detailed requirements document is in work (see development phase 1).

**Expected outcome**

Technologies developed in this task will be shown in AVL's demonstrator in WP7 / LL1, and will be applied to T7.4 "Design and validation of next generation hybrid powertrain / E-Drive". AVL will focus on addressing enhanced communication and verification requirements of mixed-criticality automotive multicore control units.

Novelty:

- Allowing multiple applications on xCU (mixed criticality) to communicate via multiple logical connections / streams through high-performance automotive networks like CAN-FD or Automotive Ethernet.

- Extend standard protocols (like CCP/XCP) synchronization of multiple data streams, addressing also QoS/RT requirements.

### 3.2.3 SevenSols (15O SevenS)

**Extending QoS and Redundancy for High-Accurate Distributed Control Systems**

Distributed control systems have strict requirements for control commands transfer (critical packets). This data should reach all network nodes in spite of network congestion (QoS) or even when any of the links between nodes experience a failure. The propagation of critical data should be deterministic and its worst case should be possible to be calculated. In addition, the network should be able to recover and reconfigure itself in case of single point of failure. The kind of applications we are focusing on related to time and phase synchronization where timing information is critical and needs to be delivered in a dependable way. This section provides some hints about the work being done by Seven Solutions on this framework.

It consists in a distributed system capable of distributing time and commands (triggering events at a specific time very accurately). In our application scenario, timing data is considered critical, whilst the rest will be considered as non-critical. White-Rabbit, known for being the most accurate time distribution solution, will be used to synchronize all network nodes. Thus, the extension of QoS, a probabilistic PI adjustment model for time phase correction and the development of redundancy protocols such as HSR (High-availability Seamless Redundancy) and PRP (Parallel Redundancy Protocol) will be focused on the propagation of White-Rabbit Ethernet frames over the network.

**Development phases:**

- Step 1: M12 (end of March 2015)
  - Concept and requirements definition
  - Redundant topologies requirements definition (SW/HW/Gateware)
- Step 2: M24 (end of March 2016)
  - Evaluation of QoS effects in terms of system jitter and wander.
  - Time accuracy tests of time-triggering using redundant network topology architectures (based on HSR).
  - First prototype of the HSR event-trigger application for distributed mixed-critical systems using White Rabbit and evaluation of QoS and VLAN support.
- Step 3: M36 (end of March 2017)
  - Time accuracy tests of time-triggering using redundant network topologies (RSTP or PRP).
  - Final EMC2 demonstrator available to whole consortium.

Requirements will be available at the end of M12 (HSR) and M24 (PRP).

**Expected outcome**

The inclusion of the redundancy protocols developed (HSR/PRP) in WP1 and integrated here with the notion of critical packets and QoS, the failure recovery system and the phase correction probabilistic algorithms in conjunction with White-Rabbit as the main time distributor, are expected to form a base for next-gen time-based systems in different domains, such as Smart Grid and Automotive.
7S will focus on the definition, development and implementation of standard protocols that enhance QoS and availability of Ethernet time-based applications in the framework of mixed-criticality systems.

Novel technology in comparison to state-of-the-art:

- It includes White-Rabbit, the most accurate PTP (IEEE-1588) solution for time and frequency distribution
- It extends standard protocols such as PTP and SyncE for time distribution for high-accurate time-triggered systems providing phase synchronization
- zero time recovery in case of failure thanks to HSR and PRP configurations
- Based on open specifications with basic open software and open Hardware reference designs.

### 3.2.4 NXP Germany (01R NXPGE)

**Multi Secure Elements Managing Unit**

NXP GE will contribute elements, e.g. architectures for multi secure element cores (multiple secure elements in a small/static system that provide security services on network requests). This requires a managing unit within this multi core to handle the external service requests and to translate them into local requests.

The objective is to develop and implement a flexible multi-core root of trust. This system shall simultaneously ensure strong security for sensitive data, fulfill increased performance requirements, and provide enhanced interoperability for other system components.

**Development phases:**

The system for a multi-secure element core will be designed in detail as a next step. The hardware development procedure will also include the implementation of corresponding security management subcomponents and firmware. Furthermore, the multi-secure element core design will be verified.

**Expected outcome**

Core component is a request handler with attached multiple secure elements as roots of trust. The request handler can be implemented on an FPGA evaluation board. The implementation will be used for the demonstrator of LL4 – in the industrial automation use case.

Offering security as a service will create the need to adopt multi-core nodes to the corresponding security requirements. In particular, security services that provide communication authenticity / integrity / confidentiality are necessary. The foundation for realizing these services is a consistent and complete chain of trust to enable end-to-end security.

Implementation of a scalable security architecture based on a security request handler to establish multiple secure communication connections between different devices based on HW security components.

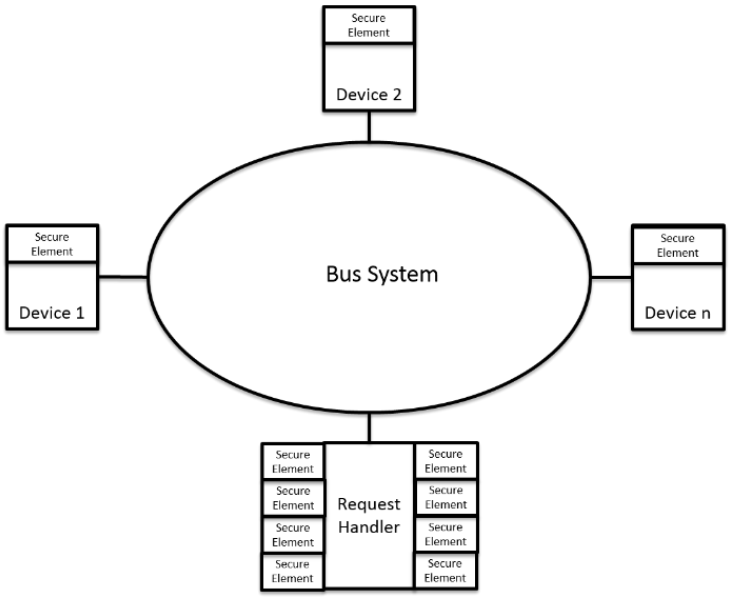


**Figure 3: Multi Secure Elements Managing Unit**

### Investigating high speed and low interference inter-chip communication and data interface preventing interference disturbed of Radio Front End.

The goal for this project is to connect two EMC ICs with a high speed link. An FPGA, which emulates a SDRAM, is chosen for that purpose. This SDRAM emulation has a second port which is used to get the data and put that over a Gigabit Ethernet connection to the other side. This is realized using a SoCrates-board produced by EBV. It contains an ALTERA FPGA Cyclone V with an embedded ARM CortexA9 subsystem, including a Gigabit Ethernet interface. The complete configuration was as follows:

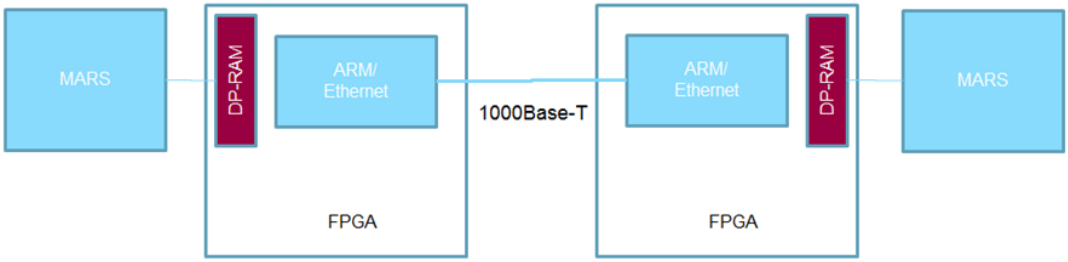


**Figure 4: High speed and low interference chip link**

### Expected outcome

Two options for the link PHY layer and protocol were investigated. A first version of the setup employed a proprietary PHY and protocol, optimized for low spectral emissions using data scrambling. As an alternative option, standard Ethernet PHY and protocol were investigated.

The standard Ethernet link appears to be well suitable, because of the signal shaping performed, and because the link is mostly silent when no data needs to be transmitted. Investigations are still going on. Special interest will be the upcoming version of automotive Gbit Ethernet that is presently standardized by the IEEE.

### 3.2.5 Infineon Technologies AG (01A IFAG)

**High speed inter-chip communication**

IFAG Munich develops a high speed serial link peripheral for inter chip communication. It provides point-to-point communication of both single data values and of large data blocks, called streams. The communicating devices can be complex microcontrollers, or a microcontroller and a device with only basic execution capabilities. There are four channels to transfer single values to/from target. They support direct writing of 8/16/32 bit data from the initiator into a target's register, as well as reading a value from a target, performed by a modules internal master on the target side. For transferring large data blocks there is a channel containing FIFOs. The module implements Transport Layer tasks and hands over the data to another module which provides Data Link Layer and Physical Layer services, data serialization and transmission. In multi slave use-cases, one master can communicate with up to three slaves.

All transfers are protected by safety features like CRC and timeout.

**Development phases**

- M1-12: Specification
- M12-24: VHDL Implementation and pre-silicon verification
- M24-36: Physical implementation, production and post-silicon verification

**Expected outcome**

The outcome of this project will be a peripheral IP module implemented on a microcontroller chip with proven functionality in a multi-slave environment. A validation and measurement report will be provided as a final delivery.

## 3.3 ST4.3.3 Peripheral components for mixed criticality systems

### 3.3.1 TU Dortmund (01W TUDO)

**Peripheral Virtualization Manager**

Heterogeneous platforms introduce additional complexities with respect to I/O addressing and arbitration. We aim to provide an abstraction layer for both abstracting as well as arbitrating concurrent accesses to I/O peripherals.

The peripheral virtualization manager will address resource sharing and arbitration on two levels:

- Address translation of core-specific memory maps will transparently enable heterogeneous platforms.
- A priority-/criticality-aware arbitration will regulate concurrent accesses to system resources.

**Development phases**

Our SystemC design-space-exploration platform will be expanded as follows:

- M1-12:
  - Platform initially supported multiple LEON3 cores. Heterogeneous cores needed for testing virtualization concepts.
  - Defined and implemented an ARM Cortex A9 ISS. Currently in integration and verification phase.

- M13-24:
  - Definition and initial implementation of peripheral virtualization manager.
  - Simulation of relevant benchmarks.
- M25-36:
  - Optimization and final verification of peripheral virtualization manager.

**Expected outcome**

We seek to provide efficient virtualization and arbitration strategies that mitigate resource sharing bottlenecks as far as possible. Furthermore, since future systems will move towards flexible and heterogeneous many-core designs, an abstraction technique will be developed that allows emulation of different virtual microcontroller systems on a single many-core while taking different criticality levels into account.

### 3.3.2 **Tecnalia (15E Tecnalia)**

**Reliable Dynamic Reconfiguration Manager**

Dynamic Partial Reconfiguration (DPR) for FPGA devices has many advantages not only for the Space domain, but also for many other industrial domains. The use of DPR requires including a module (peripheral) to control the process in a static partition of the FPGA (an area that cannot be dynamically reconfigured). Nevertheless, the space domain is a very demanding environment exposed to high radiation doses and where reliability of systems is critical. Therefore, a reliable reconfiguration manager peripheral specifically protected against radiation induced errors is highly desirable in this environment.
The peripheral will be a Microblaze based embedded system, with most of the functions related to DPR and error mitigation performed by software applications running in the Microblaze processor. The peripheral will also be able to communicate with external elements using a custom defined protocol.

Three Fault-tolerant techniques will be included for error mitigation:

- MicroBlaze fault-tolerant features
- Periodical MicroBlaze memory scrubbing done by software
- Partial bitstream data integrity check function

Work in the task will include the following activities:

- Analysis and selection of fault-tolerant techniques for embedded system design
- Design of the peripheral high-level architecture
- Design of the communications protocol
- Design of the low level architecture
- Development of the embedded system reliable hardware
- Development of the communications protocol software implementation
- Development of the dynamic partial reconfiguration management software
- Development of the memory scrubbing control software

**Expected outcome**

As mentioned before, the focus of the work is on the design of a reliable Dynamic Reconfiguration Manager Peripheral adapted to the space environment. Several error mitigation techniques will be included by design in the peripheral. This reliable peripheral will increase flexibility and scalability of the embedded system by means of DPR but guaranteeing maximum levels of service due to its error mitigation techniques. The results of the work will be shown in a specific demonstrator as described in deliverable D4.2.

The activity of the project brings two main novelties. On the one hand, the combination of DPR with radiation induced error mitigation techniques, as a way of adapting DPR to the space environment. Additionally, most of the latter techniques will be implemented by software instead of using hardware logic resources from the FPGA allowing the use of smaller devices.

# 4. Conclusions

The deliverable D4.3 contains individual contributions of partners in tasks T4.2 and T4.3. The contributions provide short an abstract of technical work performed in respective tasks and subtasks, as well as future development plans and expected results. Project activities described in this document reflect a broad spectrum of technologies implemented by multiple partners individually or in cooperation with other partners.

The task T4.2 includes topics relevant to the implementation of multi-core and many-core architectures in mixed-criticality environments. It also explores new approaches in memory and power management technologies. It also reflects on hardware reconfiguration and dynamic behavior and its correlation with mixed-criticality.

The task T4.3 explores communication paradigms on chip, between chips and between systems in context of mixed-criticality and dynamic behavior. In addition, it also investigates QoS in on- and off-chip communication channels, and novel approaches to communication with peripheral devices. Activities performed in one task frequently support or extend the work from other tasks, and contribute to complete integration. The activities in both tasks offer a wide diversity of approaches towards hardware architecture with focus on different system characteristics. The report provides an excellent overview of the partner activities in the area of multi-core hardware for mixed-criticality systems. The EMC$^2$ provides unique opportunity to showcase and extend these technologies, and to combine them with the challenges of dynamic environments. The results of the project will provide a firm foothold for the European community in above mentioned areas.

# 5. References

[1] Boris Motruk, et al., "Safe Virtual Interrupts Leveraging Distributed Shared Resources and Core-to-Core Communication on Many-Core Platforms".

[2] Boris Moturk, et al., "Power Monitoring for Mixed-Criticality on a Many-Core Platform".

[3] Boris Moturk, et al., "IDAMC: A Many-Core Platform with Run-Time Monitoring for Mixed-Criticality".

[4] UTIA, "EdkDSP," 2015. [Online]. Available: http://sp.utia.cz/index.php?ids=results&id=Utia_EdkDSP_Vivado_2013_4_EMC2..

# 6. Abbreviations

**Table 1: Abbreviations**

| Abbreviation | Meaning |
|---|---|
| µC | Micro-Controller |
| AMP | Asymmetric Multi-processing |
| AMSPS | Analog-Mixed-Signal Power System |
| AVFS | Adaptive Voltage Frequency Scaling |
| EMC$^2$ | Embedded multi-core systems for mixed criticality applications in dynamic and changeable real-time environments |
| HSR | High-availability Seamless Redundancy |
| IDAMC | Integrated Dependable Architecture for Many Cores |
| MPSoC | Multiple-processor System-on-chip |
| PRP | Parallel Redundancy Protocol |
| PTP | Precision Time Protocol |
| ToF | Time of Flight |
| TTNoC | Time Triggered Network-on-Chip |
| xCU | Electronic control unit in a car |
| DPR | Dynamic Partial Reconfiguration |
| DAL | Development Assurance Level |
| MCMC | Multi-core systems for mixed criticality applications |
| MCP | Multi-core Processor |