# ECRTS 2018

## INDUSTRIAL CHALLENGES
MOVING FROM CLASSICAL TO HIGH PERFORMANCE REAL-TIME SYSTEMS

ARNE HAMANN
BOSCH

BOSCH

# Interactive Session - Industrial Challenges
## Outline

- WATERS Challenge 2016 + 2017
  - Classical Automotive Systems Organization
  - Challenge Recap
  - Impact so far …

- WATERS Challenge 2019: Transition from µC to µP based systems
  - Trends in automotive E/E architectures
  - Analytic approaches: Evolution of performance analysis?
  - Constructive approaches: How to deal with contention effects?

# Interactive Session - Industrial Challenges
## Outline

- WATERS Challenge 2016 + 2017
  - **Classical Automotive Systems Organization**
  - Challenge Recap
  - Impact so far …
- WATERS Challenge 2019: Transition from µC to µP based systems
  - Trends in automotive E/E architectures
  - Analytic approaches: Evolution of performance analysis?
  - Constructive approaches: How to deal with contention effects?

# Interactive Session - Industrial Challenges

## Classic Automotive Software Architecture Pattern

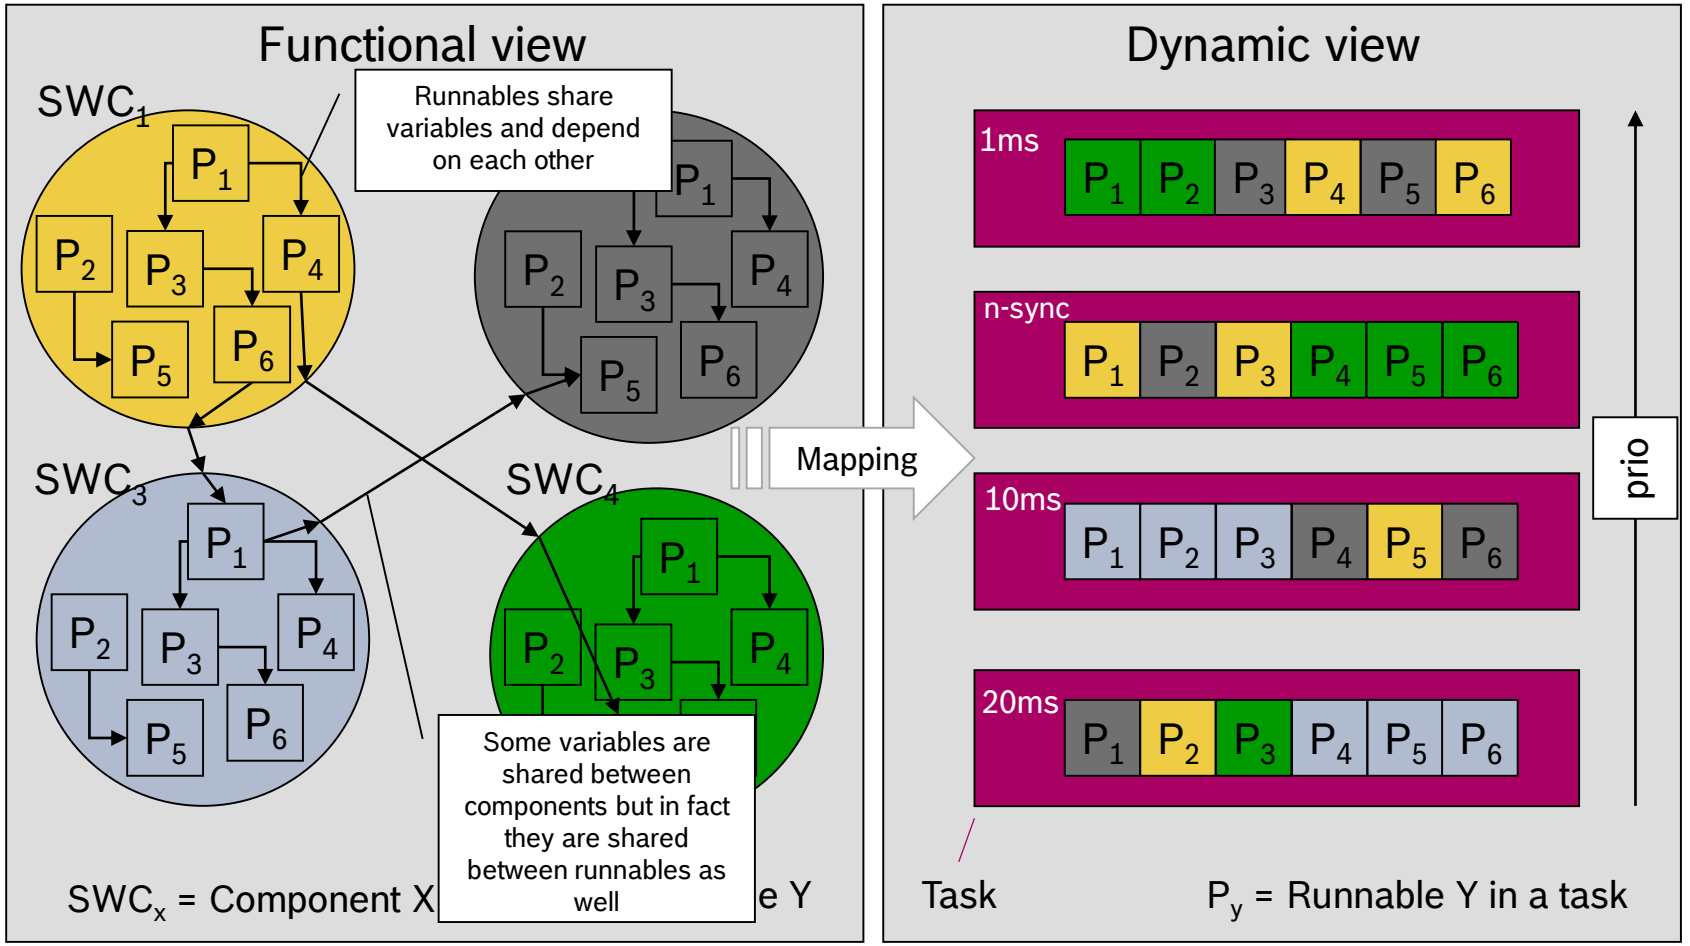

# Interactive Session - Industrial Challenges
## Outline

- WATERS Challenge 2016 + 2017
  - Classical Automotive Systems Organization
  - **Challenge Recap**
  - Impact so far …

- WATERS Challenge 2019: Transition from µC to µP based systems
  - Trends in automotive E/E architectures
  - Analytic approaches: Evolution of performance analysis?
  - Constructive approaches: How to deal with contention effects?

# Interactive Session - Industrial Challenges
## WATERS Industrial Challenges 2016/17

**https://waters2017.inria.fr/challenge/**

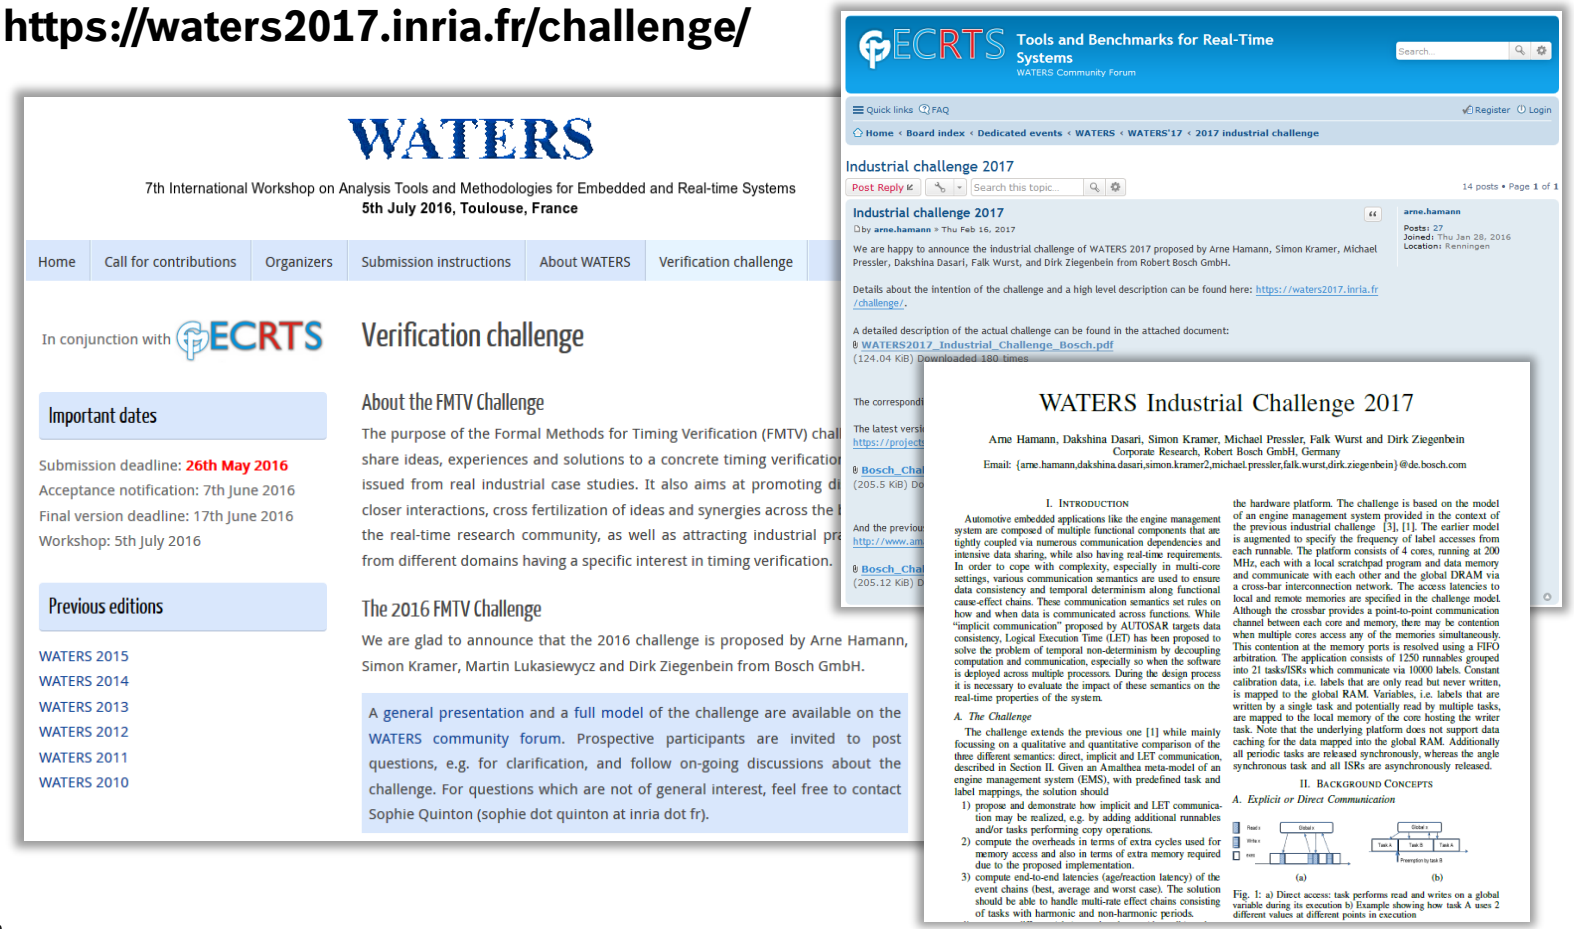

WATERS

7th International Workshop on Analysis Tools and Methodologies for Embedded and Real-time Systems
**5th July 2016, Toulouse, France**

Home | Call for contributions | Organizers | Submission instructions | About WATERS | Verification challenge

In conjunction with ECRTS

### Verification challenge

#### Important dates

Submission deadline: **26th May 2016**
Acceptance notification: 7th June 2016
Final version deadline: 17th June 2016
Workshop: 5th July 2016

#### Previous editions

WATERS 2015
WATERS 2014
WATERS 2013
WATERS 2012
WATERS 2011
WATERS 2010

#### About the FMTV Challenge

The purpose of the Formal Methods for Timing Verification (FMTV) chal
share ideas, experiences and solutions to a concrete timing verificatio
issued from real industrial case studies. It also aims at promoting di
closer interactions, cross fertilization of ideas and synergies across the
the real-time research community, as well as attracting industrial pra
from different domains having a specific interest in timing verification.

#### The 2016 FMTV Challenge

We are glad to announce that the 2016 challenge is proposed by Arne Hamann, Simon Kramer, Martin Lukasiewycz and Dirk Ziegenbein from Bosch GmbH.

A general presentation and a full model of the challenge are available on the WATERS community forum. Prospective participants are invited to post questions, e.g. for clarification, and follow on-going discussions about the challenge. For questions which are not of general interest, feel free to contact Sophie Quinton (sophie dot quinton at inria dot fr).

ECRTS Tools and Benchmarks for Real-Time Systems
WATERS Community Forum

Home ‹ Board index ‹ Dedicated events ‹ WATERS ‹ WATERS'17 ‹ 2017 industrial challenge

Industrial challenge 2017

Industrial challenge 2017
by arne.hamann » Thu Feb 16, 2017

We are happy to announce the industrial challenge of WATERS 2017 proposed by Arne Hamann, Simon Kramer, Michael Pressler, Dakshina Dasari, Falk Wurst, and Dirk Ziegenbein from Robert Bosch GmbH.

Details about the intention of the challenge and a high level description can be found here: https://waters2017.inria.fr/challenge/.

A detailed description of the actual challenge can be found in the attached document:
WATERS2017_Industrial_Challenge_Bosch.pdf
(124.04 KiB) Downloaded 180 times

WATERS Industrial Challenge 2017

Arne Hamann, Dakshina Dasari, Simon Kramer, Michael Pressler, Falk Wurst and Dirk Ziegenbein
Corporate Research, Robert Bosch GmbH, Germany
Email: {arne.hamann,dakshina.dasari,simon.kramer2,michael.pressler,falk.wurst,dirk.ziegenbein}@de.bosch.com

I. Introduction

Automotive embedded applications like the engine management system are composed of multiple functional components that are tightly coupled via numerous communication dependencies and intensive data sharing, while also having real-time requirements. In order to cope with complexity, especially in multi-core settings, various communication semantics are used to ensure data consistency and temporal determinism along functional cause-effect chains. These communication semantics set rules on how and when data is communicated across functions. While "implicit communication" proposed by AUTOSAR targets data consistency, Logical Execution Time (LET) has been proposed to solve the problem of temporal non-determinism by decoupling computation and communication, especially so when the software is deployed across multiple processors. During the design process it is necessary to evaluate the impact of these semantics on the real-time properties of the system.

A. The Challenge

The challenge extends the previous one [1] while mainly focussing on a qualitative and quantitative comparison of the three different semantics: direct, implicit and LET communication, described in Section II. Given an Amalthea meta-model of an engine management system (EMS), with predefined task and label mappings, the solution should

1) propose and demonstrate how implicit and LET communication may be realized, e.g. by adding additional runnables and/or tasks performing copy operations.
2) compute the overheads in terms of extra cycles used for memory access and also in terms of extra memory required due to the proposed implementation.
3) compute end-to-end latencies (age/reaction latency) of the event chains (best, average and worst case). The solution should be able to handle multi-rate effect chains consisting of tasks with harmonic and non-harmonic periods.

the hardware platform. The challenge is based on the model of an engine management system provided in the context of the previous industrial challenge [3], [1]. The earlier model is augmented to specify the frequency of label accesses from each runnable. The platform consists of 4 cores, running at 200 MHz, each with a local scratchpad program and data memory and communicate with each other and the global DRAM via a cross-bar interconnection network. The access latencies to local and remote memories are specified in the challenge model. Although the crossbar provides a point-to-point communication channel between each core and memory, there may be contention when multiple cores access any of the memories simultaneously. This contention at the memory ports is resolved using a FIFO arbitration. The application consists of 1250 runnables grouped into 21 tasks/ISRs which communicate via 10000 labels. Constant calibration data, i.e. labels that are only read but never written, is mapped to the global RAM. Variables, i.e. labels that are written by a single task and potentially read by multiple tasks, are mapped to the local memory of the core hosting the writer task. Note that the underlying platform does not support data caching for the data mapped into the global RAM. Additionally all periodic tasks are released synchronously, whereas the angle synchronous task and all ISRs are asynchronously released.

II. Background Concepts

A. Explicit or Direct Communication




Fig. 1: a) Direct access: task performs read and writes on a global variable during its execution b) Example showing how task A uses 2 different values at different points in execution

# Interactive Session - Industrial Challenges
## Challenge HW Model

APP4MC

- Simplified AURIX Architecture

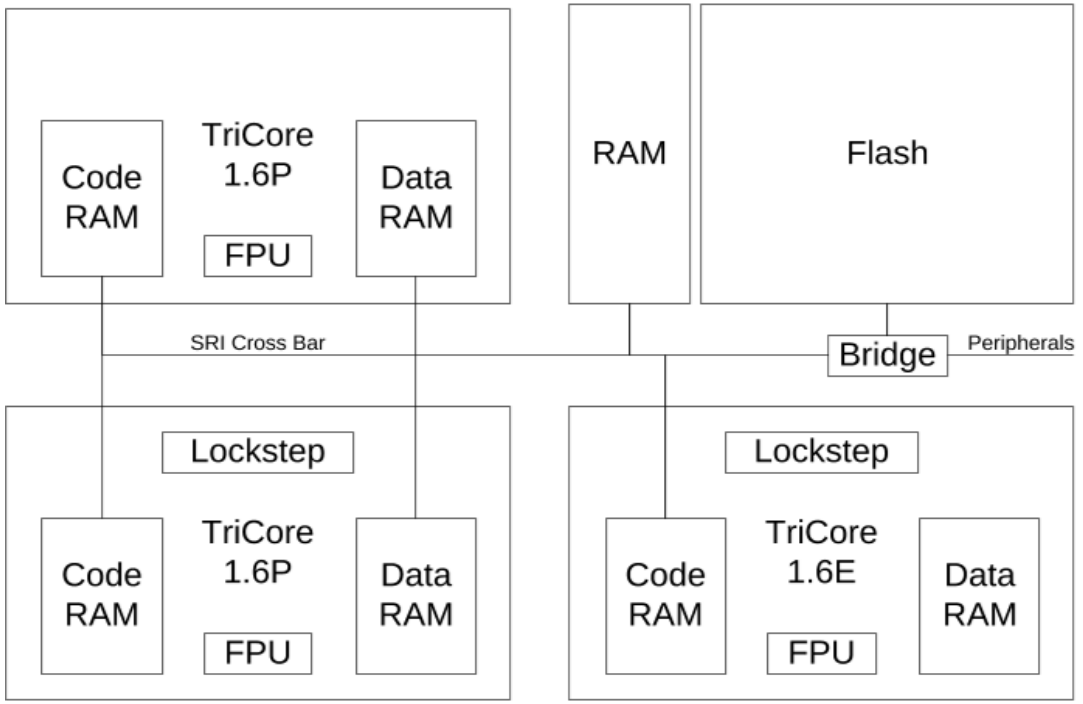


- Memory Access Time

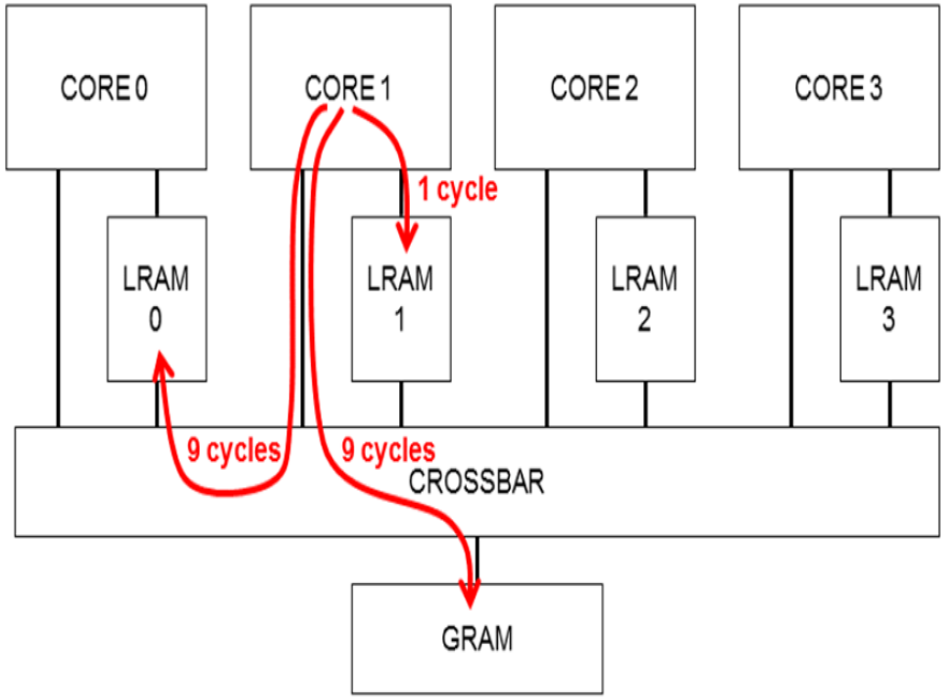





BOSCH

# Interactive Session - Industrial Challenges
## Challenge SW Model


- Key data of the model
  - **1250 Runnables** mapped to
  - **21 Tasks & Interrupts** accessing
  - **10.000 Labels** (shared data)
  - Event chains

- Huge amount of data dependencies
  - challenge exact analysis methods

| I | II | III | IV | V | VI |
|---|---|---|---|---|---|
| <10 | 10-50 | 51-100 | 100-500 | 501-1000 | >1000 |

TABLE II. INTER-TASK COMMUNICATION

| Period | 1 ms | 2 ms | 5 ms | 10 ms | 20 ms | 50 ms | 100 ms | 200 ms | 1000 ms | sync |
|---|---|---|---|---|---|---|---|---|---|---|
| 1 ms | | | | I | I | | I | | | I |
| 2 ms | | | | I | I | | I | | | |
| 5 ms | | I | IV | IV | II | II | I | | | |
| 10 ms | II | II | II | VI | IV | II | IV | II | III | IV |
| 20 ms | I | I | I | IV | VI | II | IV | I | II | IV |
| 50 ms | | | II | II | II | III | I | | | |
| 100 ms | | I | I | V | IV | II | VI | II | III | IV |
| 200 ms | | | | I | I | | I | I | I | |
| 1000 ms | | | | III | II | | III | I | IV | I |
| Angle-sync | I | I | I | IV | IV | I | III | I | I | V |

# Interactive Session - Industrial Challenges

## Handle the Real Complexity…

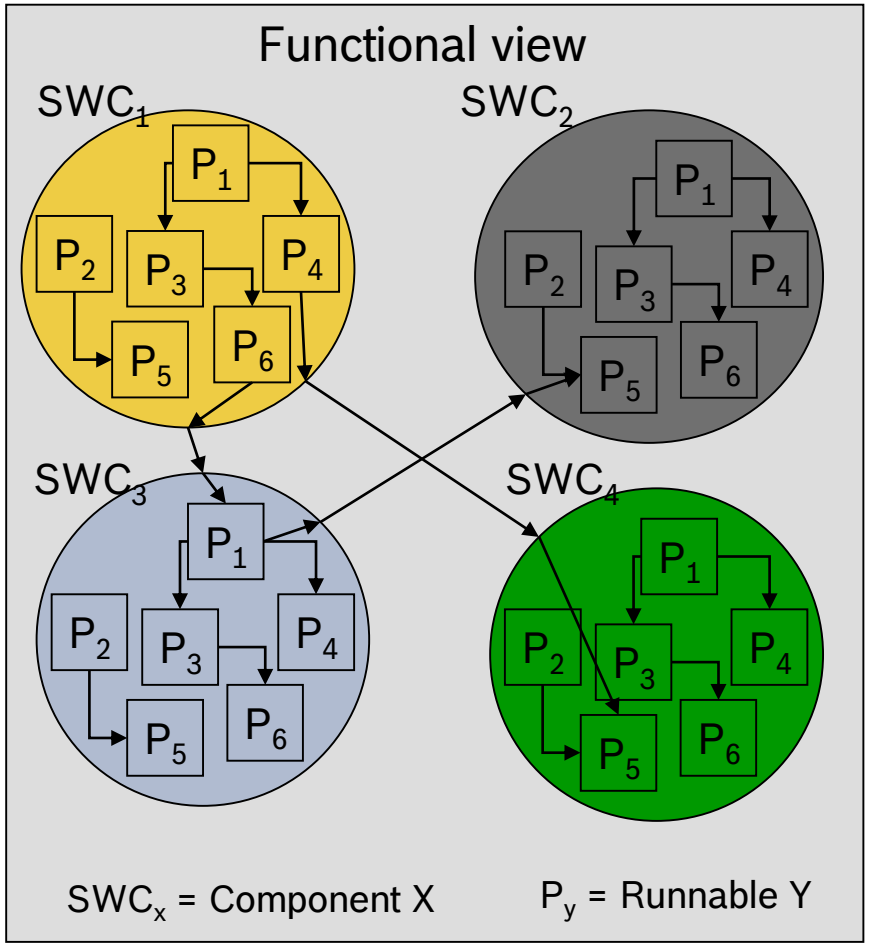


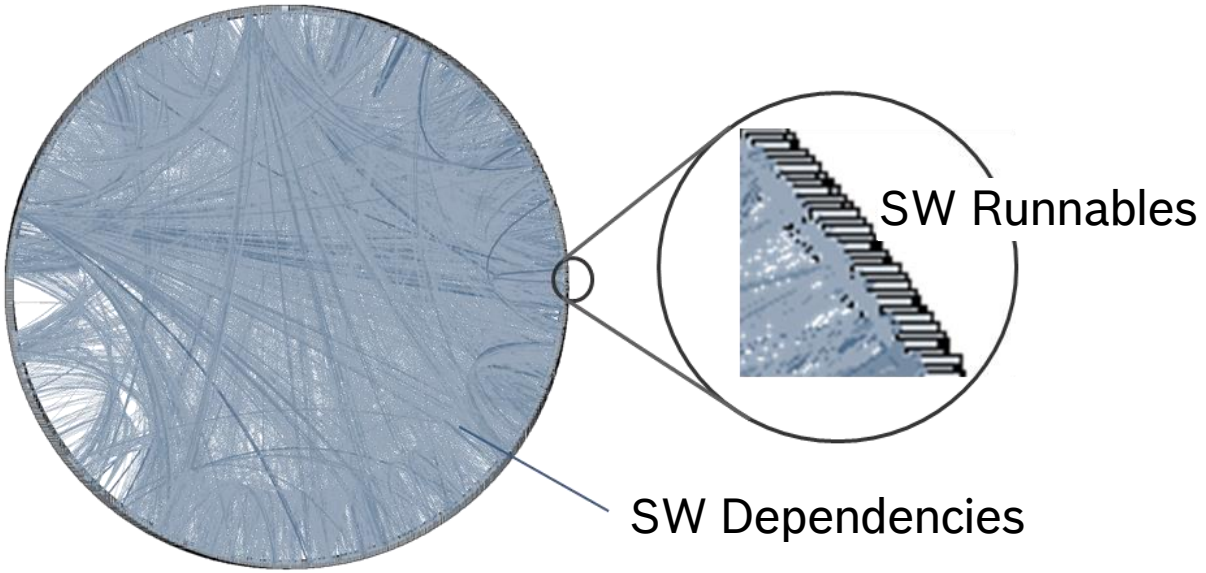


Fine-grain, legacy SW sharing between OEM and Tier1 with multiple dependencies

# Interactive Session - Industrial Challenges
## Challenge questions 2016

- **2016: focus on impact of memory contention on execution times and scheduling**
  - Calculate tight end-to-end latencies ignoring memory accesses and arbitration
  - Calculate tight end-to-end latencies **including memory access and arbitration**
  - Optimize end-to-end latencies by mapping the labels among the local and global memories

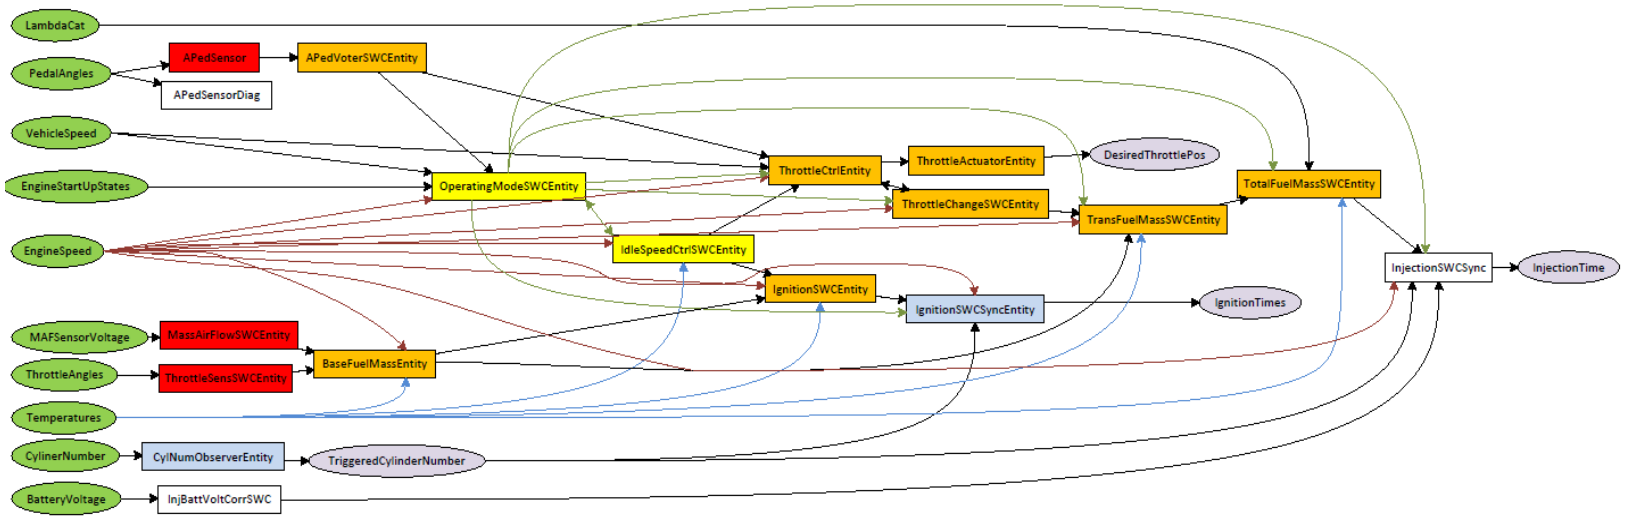


**Benchmarking, System Design and Case-studies for Multi-core based Embedded Automotive Systems**
Piotr Dziurzanski, Amit Kumar Singh, Leandro S. Indrusiak, Björn Saballus



BOSCH

# Interactive Session - Industrial Challenges
## Challenge questions 2017

- **2017: focus on impact of communication mechanisms on end-to-end latencies**
  - Proposing concepts for realizing implicit and LET communication
  - Compute the overhead induced by the proposed realizations
  - Compute the end-to-end latency along event chains due to the proposed realizations
  - Propose alternative label mapping to reduce memory access overheads
  - Consider effects of memory contention on the end-to-end latencies and memory access time

- Implicit Communication
  - Goal: data consistency

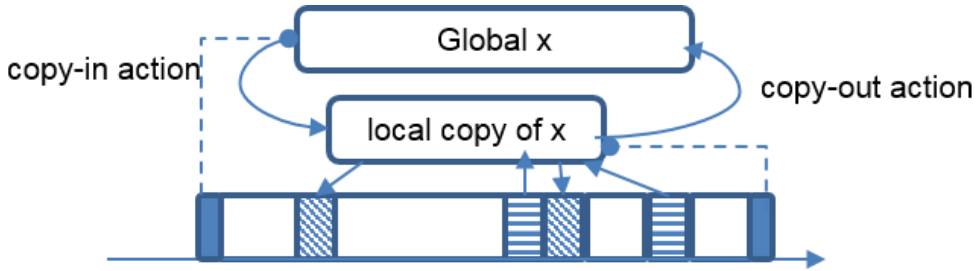


- LET Communication
  - Goal: determinism

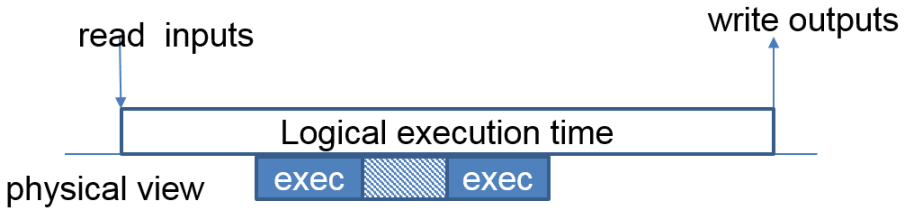

# Interactive Session - Industrial Challenges
## Outline

- WATERS Challenge 2016 + 2017
  - Classical Automotive Systems Organization
  - Challenge Recap
  - **Impact so far …**

- WATERS Challenge 2019: Transition from µC to µP based systems
  - Trends in automotive E/E architectures
  - Analytic approaches: Evolution of performance analysis?
  - Constructive approaches: How to deal with contention effects?

# Interactive Session - Industrial Challenges

## Impact on the Forum

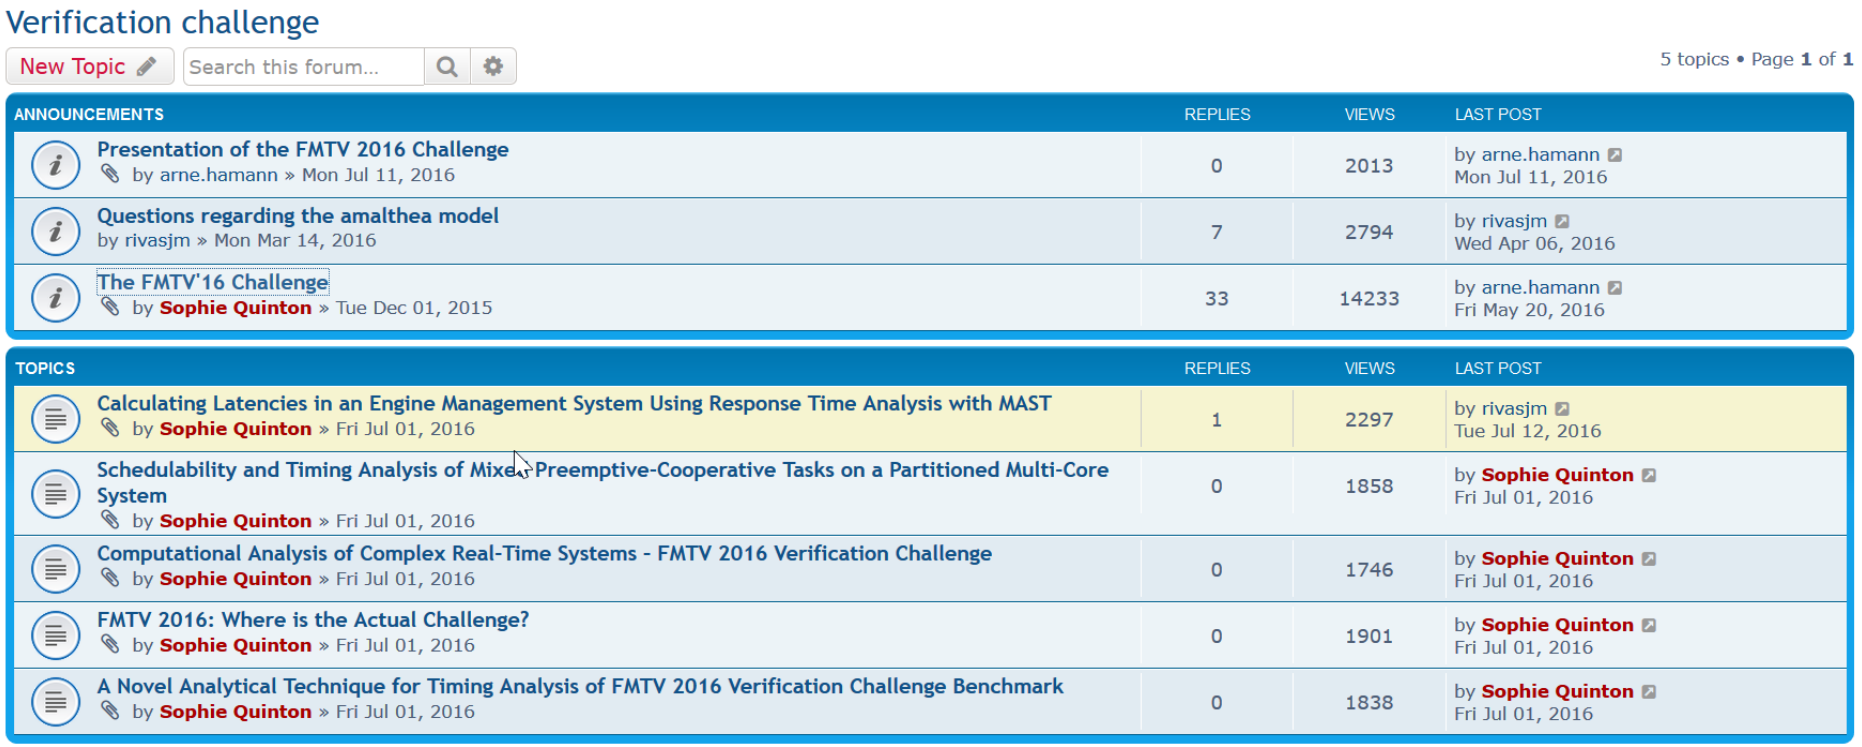

Verification challenge

New Topic | Search this forum...

5 topics • Page 1 of 1

| ANNOUNCEMENTS | REPLIES | VIEWS | LAST POST |
|---|---|---|---|
| Presentation of the FMTV 2016 Challenge<br>by arne.hamann » Mon Jul 11, 2016 | 0 | 2013 | by arne.hamann<br>Mon Jul 11, 2016 |
| Questions regarding the amalthea model<br>by rivasjm » Mon Mar 14, 2016 | 7 | 2794 | by rivasjm<br>Wed Apr 06, 2016 |
| The FMTV'16 Challenge<br>by **Sophie Quinton** » Tue Dec 01, 2015 | 33 | 14233 | by arne.hamann<br>Fri May 20, 2016 |

| TOPICS | REPLIES | VIEWS | LAST POST |
|---|---|---|---|
| Calculating Latencies in an Engine Management System Using Response Time Analysis with MAST<br>by **Sophie Quinton** » Fri Jul 01, 2016 | 1 | 2297 | by rivasjm<br>Tue Jul 12, 2016 |
| Schedulability and Timing Analysis of Mixed-Preemptive-Cooperative Tasks on a Partitioned Multi-Core System<br>by **Sophie Quinton** » Fri Jul 01, 2016 | 0 | 1858 | by **Sophie Quinton**<br>Fri Jul 01, 2016 |
| Computational Analysis of Complex Real-Time Systems - FMTV 2016 Verification Challenge<br>by **Sophie Quinton** » Fri Jul 01, 2016 | 0 | 1746 | by **Sophie Quinton**<br>Fri Jul 01, 2016 |
| FMTV 2016: Where is the Actual Challenge?<br>by **Sophie Quinton** » Fri Jul 01, 2016 | 0 | 1901 | by **Sophie Quinton**<br>Fri Jul 01, 2016 |
| A Novel Analytical Technique for Timing Analysis of FMTV 2016 Verification Challenge Benchmark<br>by **Sophie Quinton** » Fri Jul 01, 2016 | 0 | 1838 | by **Sophie Quinton**<br>Fri Jul 01, 2016 |

# Interactive Session - Industrial Challenges

## Impact on the Forum

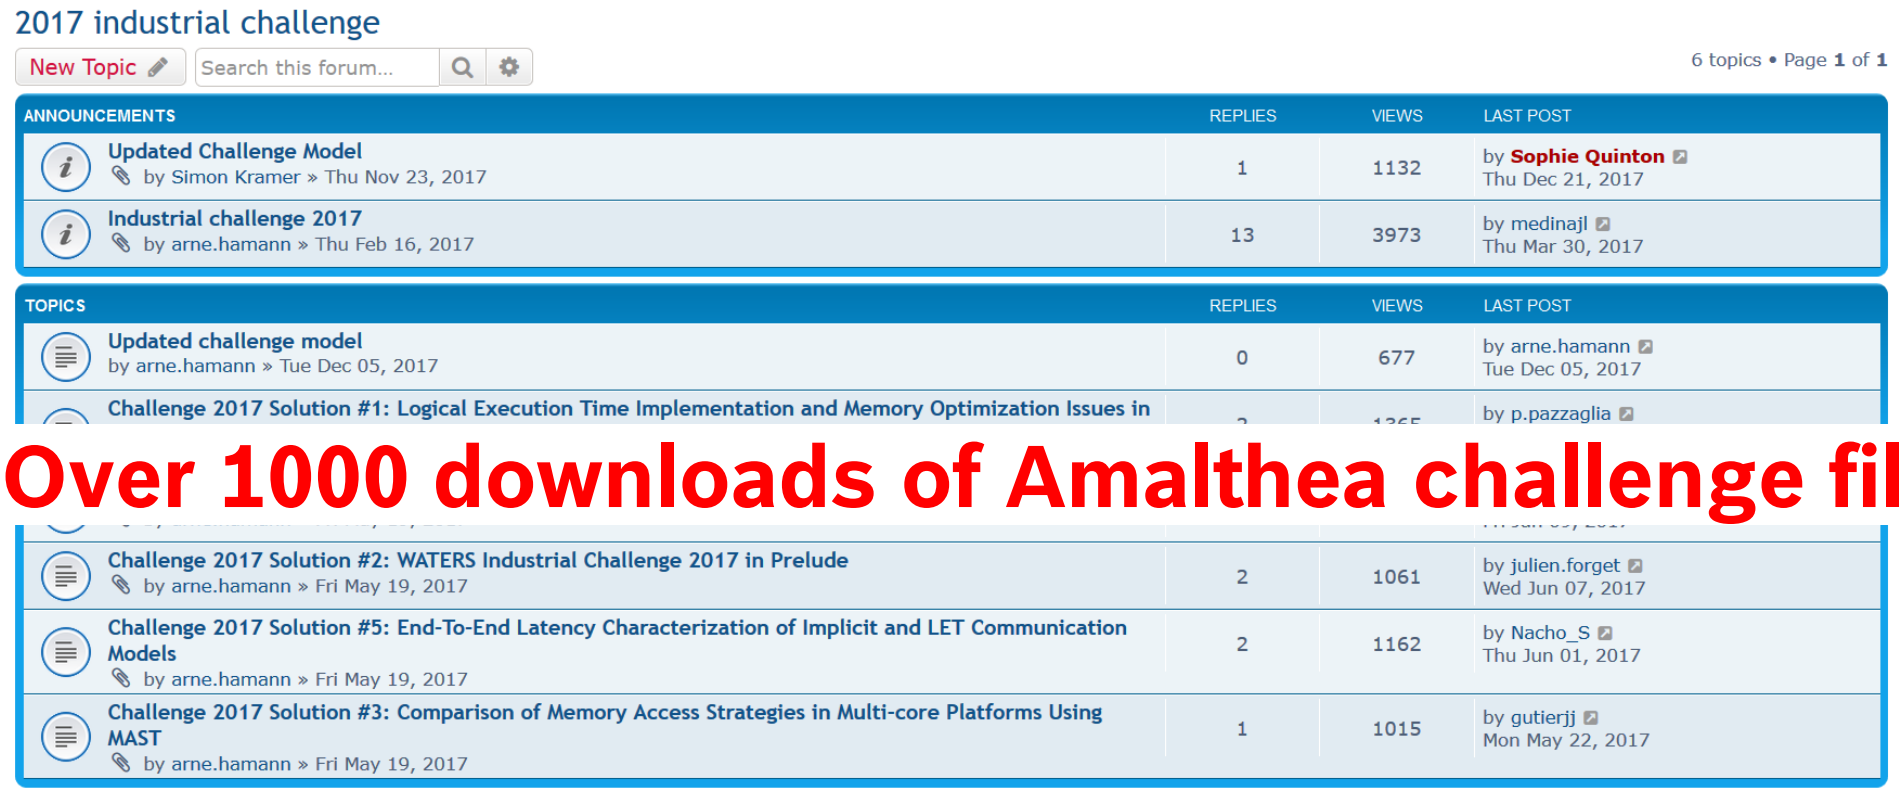

2017 industrial challenge

New Topic | Search this forum... | 6 topics • Page 1 of 1

| ANNOUNCEMENTS | REPLIES | VIEWS | LAST POST |
|---|---|---|---|
| Updated Challenge Model<br>by Simon Kramer » Thu Nov 23, 2017 | 1 | 1132 | by Sophie Quinton<br>Thu Dec 21, 2017 |
| Industrial challenge 2017<br>by arne.hamann » Thu Feb 16, 2017 | 13 | 3973 | by medinajl<br>Thu Mar 30, 2017 |

| TOPICS | REPLIES | VIEWS | LAST POST |
|---|---|---|---|
| Updated challenge model<br>by arne.hamann » Tue Dec 05, 2017 | 0 | 677 | by arne.hamann<br>Tue Dec 05, 2017 |
| Challenge 2017 Solution #1: Logical Execution Time Implementation and Memory Optimization Issues in | | | by p.pazzaglia |
| Challenge 2017 Solution #2: WATERS Industrial Challenge 2017 in Prelude<br>by arne.hamann » Fri May 19, 2017 | 2 | 1061 | by julien.forget<br>Wed Jun 07, 2017 |
| Challenge 2017 Solution #5: End-To-End Latency Characterization of Implicit and LET Communication Models<br>by arne.hamann » Fri May 19, 2017 | 2 | 1162 | by Nacho_S<br>Thu Jun 01, 2017 |
| Challenge 2017 Solution #3: Comparison of Memory Access Strategies in Multi-core Platforms Using MAST<br>by arne.hamann » Fri May 19, 2017 | 1 | 1015 | by gutierjj<br>Mon May 22, 2017 |

**Over 1000 downloads of Amalthea challenge file**

# Interactive Session - Industrial Challenges

## Impact Citations

- 58 citations of base paper (source Google Scholar)
- “Dark figure” probably higher

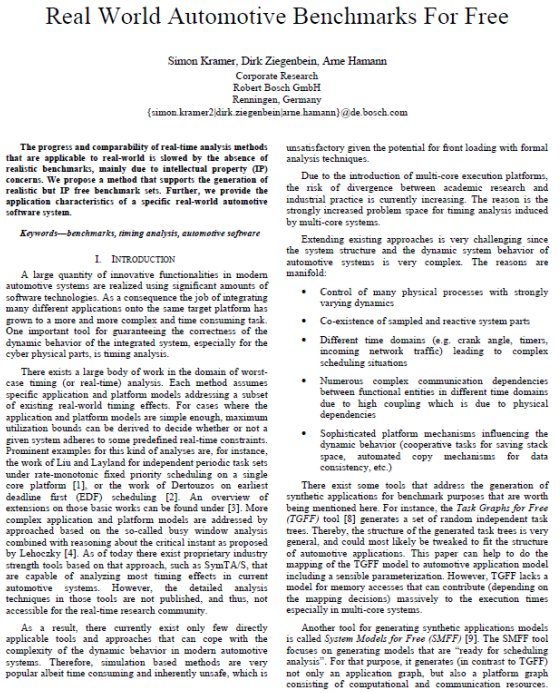

Real World Automotive Benchmarks For Free

Simon Kramer, Dirk Ziegenbein, Arne Hamann
Corporate Research
Robert Bosch GmbH
Renningen, Germany
{simon.kramer}{dirk.ziegenbein|arne.hamann}@de.bosch.com

**The progress and comparability of real-time analysis methods that are applicable to real-world is slowed by the absence of realistic benchmarks, mainly due to intellectual property (IP) concerns. We propose a method that supports the generation of realistic but IP free benchmark sets. Further, we provide the application characteristics of a specific real-world automotive software system.**

***Keywords—benchmarks, timing analysis, automotive software***

I. INTRODUCTION

A large quantity of innovative functionalities in modern automotive systems are realized using significant amounts of software technologies. As a consequence the job of integrating many different applications onto the same target platform has grown to a more and more complex and time consuming task. One important tool for guaranteeing the correctness of the dynamic behavior of the integrated system, especially for the cyber physical parts, is timing analysis.

There exists a large body of work in the domain of worst-case timing (or real-time) analysis. Each method assumes specific application and platform models addressing a subset of existing real-world timing effects. For cases where the application and platform models are simple enough, maximum utilization bounds can be derived to decide whether or not a given system adheres to some predefined real-time constraints. Prominent examples for this kind of analyses are, for instance, the work of Liu and Layland for independent periodic task sets under rate-monotonic fixed priority scheduling on a single core platform [1], or the work of Dertouzos on earliest deadline first (EDF) scheduling [2]. An overview of extensions on those basic works can be found under [3]. More complex application and platform models are addressed by approached based on the so-called busy window analysis combined with reasoning about the critical instant as proposed by Lehoczky [4]. As of today there exist proprietary industry strength tools based on that approach, such as SymTA/S, that are capable of analyzing most timing effects in current automotive systems. However, the detailed analysis techniques in those tools are not published, and thus, not accessible for the real-time research community.

As a result, there currently exist only few directly applicable tools and approaches that can cope with the complexity of the dynamic behavior in modern automotive systems. Therefore, simulation based methods are very popular albeit time consuming and inherently unsafe, which is unsatisfactory given the potential for front loading with formal analysis techniques.

Due to the introduction of multi-core execution platforms, the risk of divergence between academic research and industrial practice is currently increasing. The reason is the strongly increased problem space for timing analysis induced by multi-core systems.

Extending existing approaches is very challenging since the system structure and the dynamic system behavior of automotive systems is very complex. The reasons are manifold:

- Control of many physical processes with strongly varying dynamics
- Co-existence of sampled and reactive system parts
- Different time domains (e.g. crank angle, timers, incoming network traffic) leading to complex scheduling situations
- Numerous complex communication dependencies between functional entities in different time domains due to high coupling which is due to physical dependencies
- Sophisticated platform mechanisms influencing the dynamic behavior (cooperative tasks for saving stack space, automated copy mechanisms for data consistency, etc.)

There exist some tools that address the generation of synthetic applications for benchmark purposes that are worth being mentioned here. For instance, the *Task Graphs for Free (TGFF)* tool [8] generates a set of random independent task trees. Thereby, the structure of the generated task trees is very general, and could most likely be tweaked to fit the structure of automotive applications. This paper can help to do the mapping of the TGFF model to automotive application model including a sensible parameterization. However, TGFF lacks a model for memory accesses that can contribute (depending on the mapping decisions) massively to the execution times especially in multi-core systems.

Another tool for generating synthetic applications models is called *System Modelz for Free (SMFF)* [9]. The SMFF tool focuses on generating models that are “ready for scheduling analysis”. For that purpose, it generates (in contrast to TGFF) not only an application graph, but also a platform graph consisting of computational and communication resources.

### Real world automotive benchmarks for free

| | |
|---|---|
| Autoren | Simon Kramer, Dirk Ziegenbein, Arne Hamann |
| Publikationsdatum | 2015/7/7 |
| Zeitschrift | 6th International Workshop on Analysis Tools and Methodologies for Embedded and Real-time Systems (WATERS) |
| Zitate insgesamt | Zitiert von: 58 |

# Interactive Session - Industrial Challenges
## Outline

- WATERS Challenge 2016 + 2017
  - Classical automotive systems organization
  - Challenge recap
  - Impact so far …

- WATERS Challenge 2019: Transition from µC to µP based systems
  - **Trends in automotive E/E architectures**
  - Analytic approaches: Evolution of performance analysis?
  - Constructive approaches: How to deal with contention effects?

# Interactive Session - Industrial Challenges

## Trends in automotive E/E systems

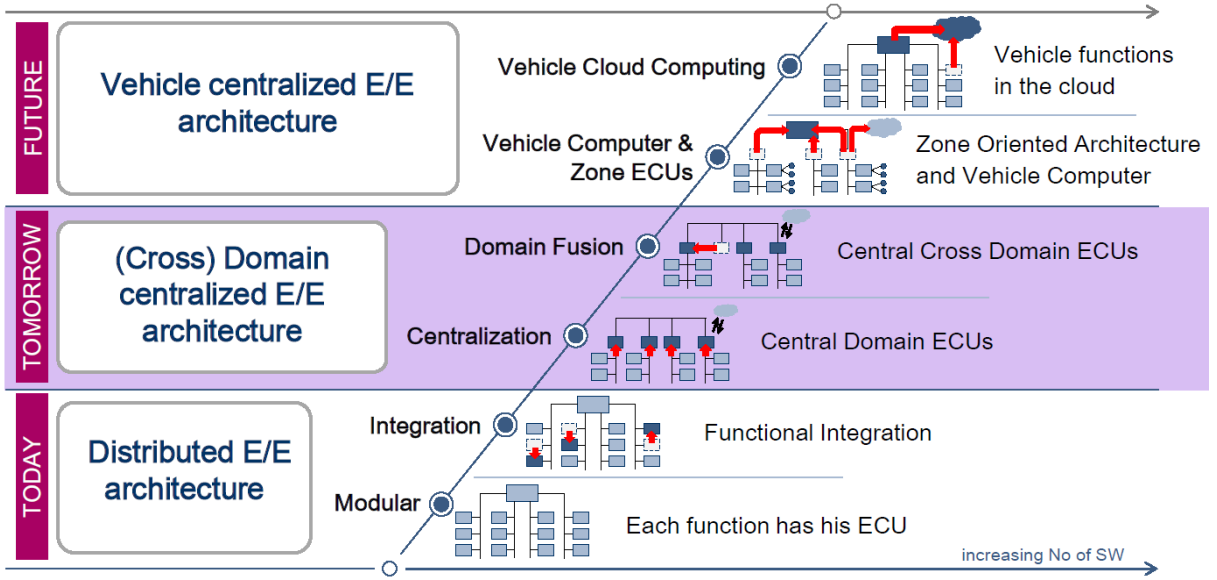


*Source: Bosch*

**Large-scale integration of heterogeneous applications** on (Cross)-Domain & Vehicle Centralized E/E Architectures

**Computing Power Demand**

Serial computing in embedded systems is hitting the **techno-logical limits**

2025 ⇧ 2015 **Factor 2 - 20**

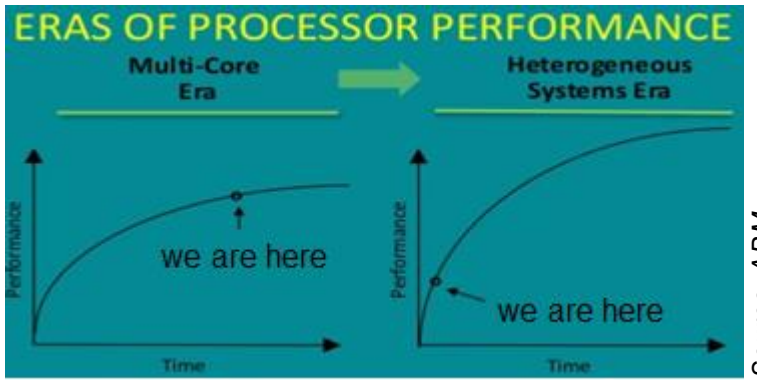


*Source: ARM*

**Heterogeneous HW platforms** to satisfy tremendous need for computing power




BOSCH

# Interactive Session - Industrial Challenges
## Outline

- WATERS Challenge 2016 + 2017
  - Classical automotive systems organization
  - Challenge recap
  - Impact so far …

- WATERS Challenge 2019: Transition from µC to µP based systems
  - Trends in automotive E/E architectures
  - **Analytic approaches: Evolution of performance analysis?**
  - Constructive approaches: How to deal with contention effects?

# Interactive Session - Industrial Challenges

## Status Quo of Performance Analysis @ Bosch

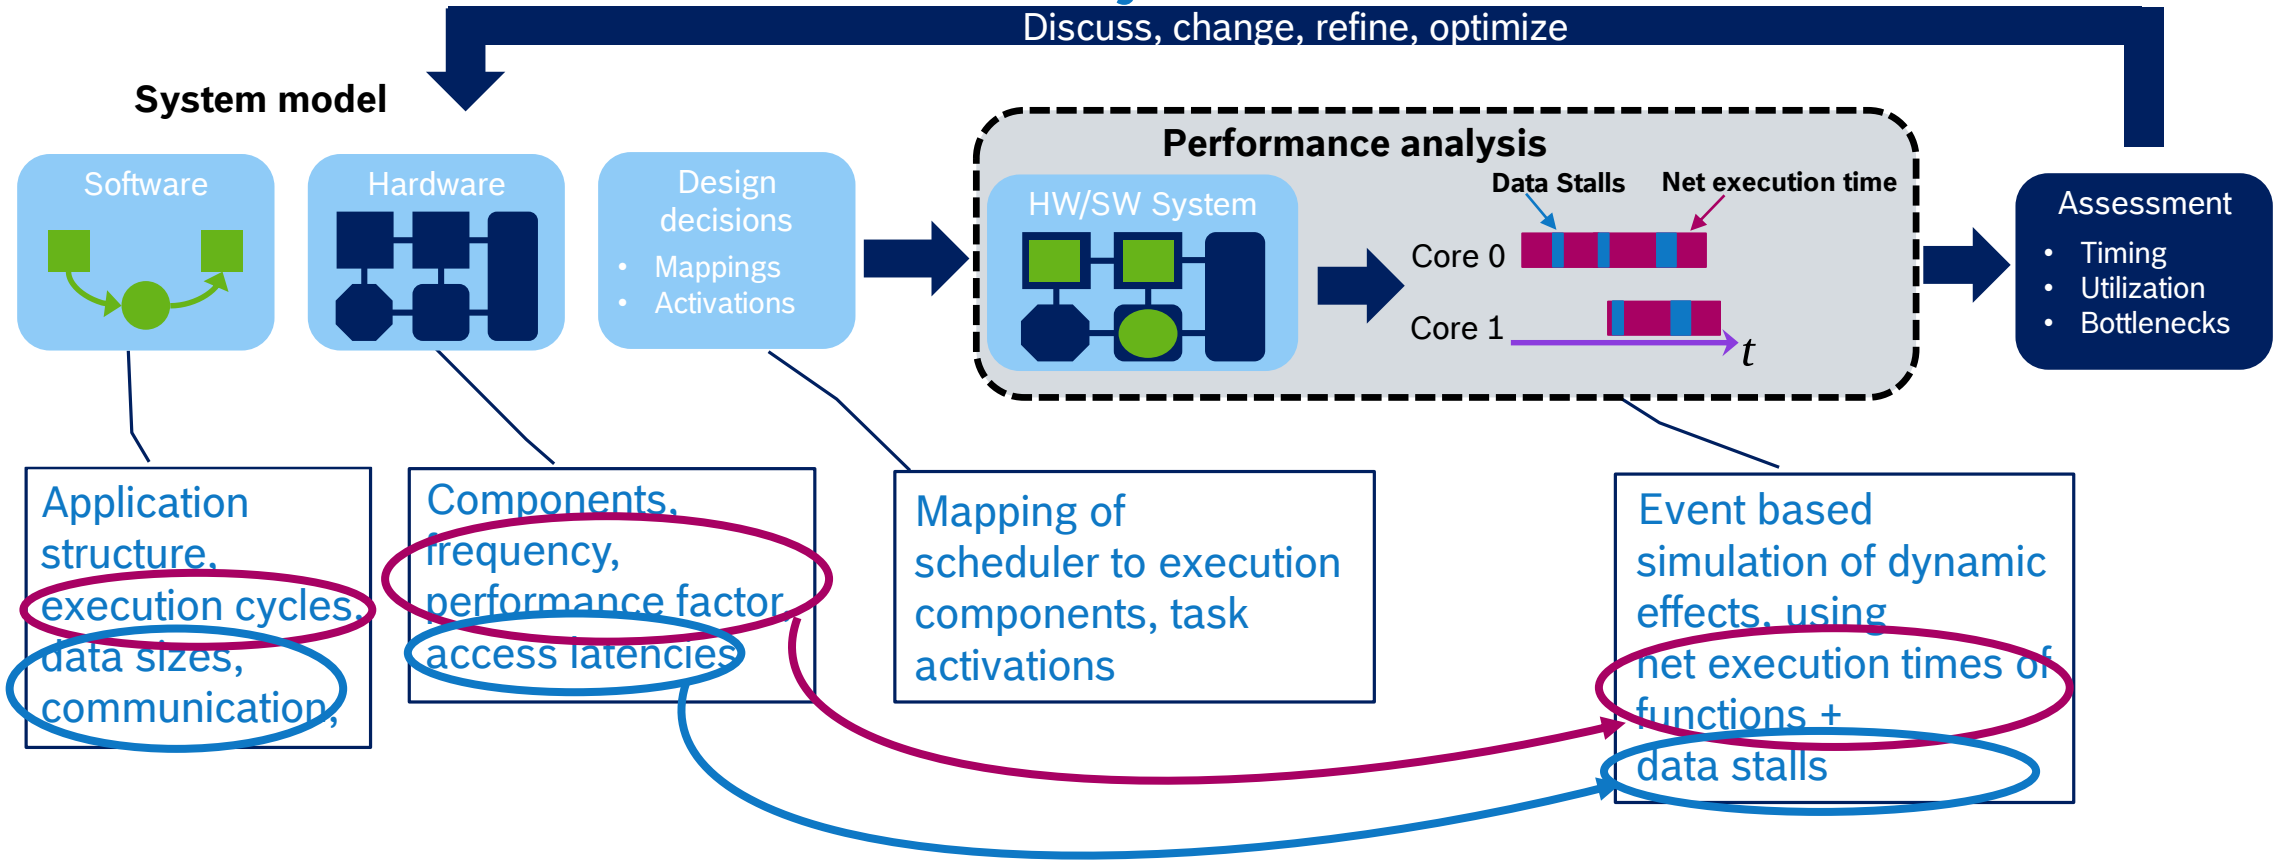


**Current modelling approach suited for (homogeneous) many-core architectures**

# Interactive Session - Industrial Challenges

## Evolution of Performance Analysis to Heterogeneous Platforms

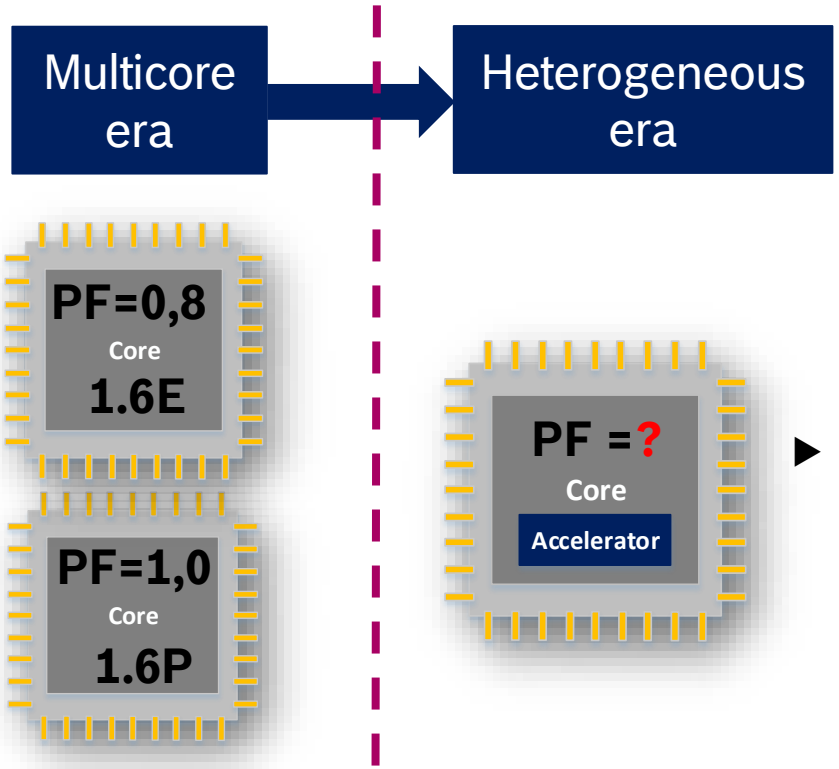


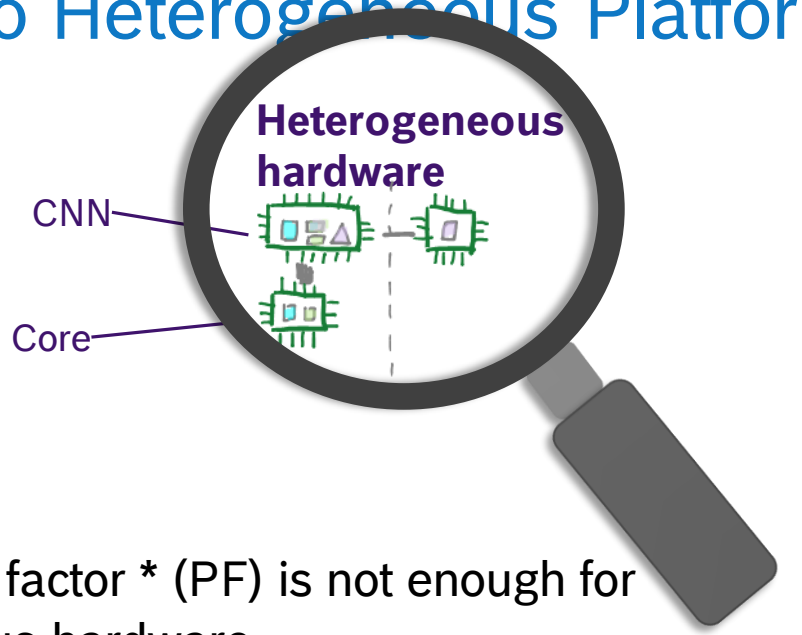


- Performance factor * (PF) is not enough for heterogeneous hardware
  - Non linear performance effects due to specific acceleration
  - PF is not transparent regarding heterogeneous effects
  - Infeasible to compare different ISAs

* E.g., IPC (Instructions Per Cycle)

**Goal: Enable models & simulation for heterogeneous software and hardware**

# Interactive Session - Industrial Challenges
## Research Challenges for Analytic Approaches

- Can we find analytic models with adequate precision ?

- A new model for execution times
  - Instruction mixes?
  - How to model the (heterogeneous) execution platform?
    - Execution platform modelling virtually not present in the real-time community

- How to predict performance of software on heterogeneous HW platforms?

- How to analytically capture the impact of communication/memory accesses?

# Interactive Session - Industrial Challenges
## Outline

- WATERS Challenge 2016 + 2017
  - Classical automotive systems organization
  - Challenge recap
  - Impact so far …

- WATERS Challenge 2019: Transition from µC to µP based systems
  - Trends in automotive E/E architectures
  - Analytic approaches: Evolution of performance analysis?
  - **Constructive approaches: How to deal with contention effects?**

# Interactive Session - Industrial Challenges
## Predictability on High-Performance Platforms

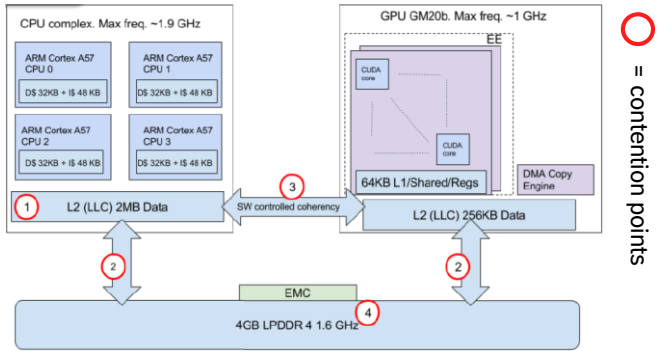


*NVIDIA Tegra X1 Platform*

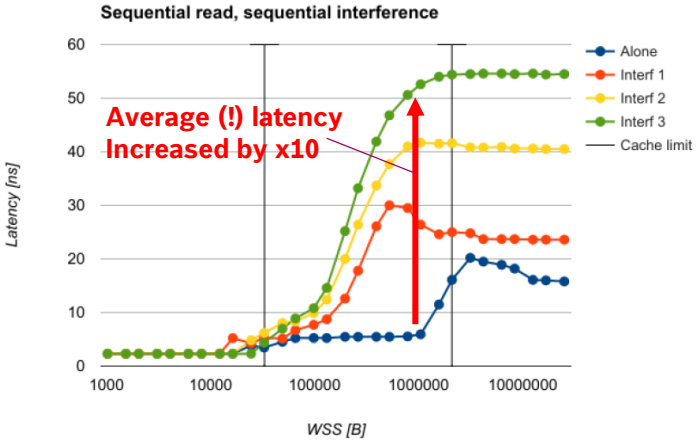


*Avg. memory access latencies per word*

Source: Roberto Cavicchioli, Nicola Capodieci, Marko Bertogna, Memory interference characterization between CPU cores and integrated GPUs in mixed-criticality platforms. ETFA 2017

- Shared memory is a big bottleneck in high-end µP based real-time platforms
- **Interference effects are more severe by orders of magnitude compared to µC platforms**
- Goal for automotive systems engineering: predictable real-time behavior on high-performance platforms



BOSCH

# Interactive Session - Industrial Challenges
## Need for Constructive Approaches

- Goal: reconcile high-performance computing with real-time predictability and determinism
- Reduce complexity and pessimism of analytic approaches
- Several groups are working on constructive mechanisms to “regulate” contention effects

**Memory Bandwidth Management for Efficient Performance Isolation in Multi-core Platforms**
Heechul Yun, Gang Yao, Rodolfo Pellizzoni, Marco Caccamo and Lui Sha
IEEE Transactions on Computers, February 2016, Volume 65(2): 562-576
**Schedulability Analysis for Memory Bandwidth Regulated Multicore Real-Time Systems**
Gang Yao, Heechul Yun, Zheng Pei Wu, Rodolfo Pellizzoni, Marco Caccamo, Lui Sha
IEEE Transactions on Computers, February 2016, Volume 65(2): 601-614
**Contention-Aware Dynamic Memory Bandwidth Isolation with Predictability in COTS Multicores: An Avionics Case Study**
Ankit Agrawal, Gerhard Fohler, Johannes Freitag, Jan Nowotsch, Sascha Uhrig, Michael Paulitsch
29th Euromicro Conference on Real-Time Systems (ECRTS), June 2017.
**SiGAMMA: Server based integrated GPU Arbitration Mechanism for Memory Accesses**
Nicola Capodieci, Roberto Cavicchioli, Paolo Valente and Marko Bertogna
Proceedings of the 25th International Conference on Real-Time Networks and Systems (RTNS’17), Grenoble, France, October 2017.me Systems (ECRTS), June 2017.

# Interactive Session - Industrial Challenges
## Bosch Challenge 2019

- In cooperation with UNIMORE, HiPeRT Lab

- Amalthea model comprising
  - Heterogeneous HW platform (µPs including accelerators such as GPU)
  - Heterogeneous SW mix (ADAS applications)

- Executable software for relevant COTS HW platform (to be defined)

- Challenges
  - Analyze response times / contention effects, including comparison with real execution times
  - Explore the impact of constructive platform approaches for the gives setup
    - Compare different approaches
    - Quantify effects for realistic automotive applications

# THANK YOU

...AND HARALD MACKAMUL, JÖRG TESSMER, FALK WURST, TOBIAS BEICHTER, SYED AOUN RAZA, DIRK ZIEGENBEIN, JENS GLADIGAU, MICHAEL PRESSLER, DAKSHINA DASARI

BOSCH