IST-214373 ArtistDesign
Network of Excellence
on Design for Embedded Systems

Activity Progress Report for Year 3

Platform and MPSoC Analysis

Cluster:

Hardware Platforms and MPSoC

Activity Leader:

Prof. Jan Madsen (DTU)
http://www.imm.dtu.dk

Policy Objective (abstract)

The main objective of the activity is to build a common research environment, which integrates performance analysis algorithms and tools for hardware platforms and Multi-Processor System-on-Chip (MPSoC). The main challenge is the introduction of new aspects such as robustness, adaptivity and power consumption, which need to be addressed at run-time. The teams involved in the activity aim at developing and integrating modeling and analysis techniques for scalable performance analysis of applications executing on embedded hardware platforms.

Versions

number	comment	date
1.0	First version delivered to the reviewers	February 4th 2010

Table of Contents

1. Overview of the Activity

1.1 ArtistDesign participants and their role within the Activity

Activity leader: Prof. Jan Madsen – Technical University of Denmark, DTU (Denmark)
System-level modeling and analysis of MPSoC and networked embedded systems. Architectures and programming models for multi-core embedded systems. Analysis and Optimization of real-time and fault-tolerant applications implemented on distributed Platforms and MPSoC. Reconfigurable platforms and run-time resource management.

Team leader: Prof. Petru Eles – Linköping University, LiU (Sweden)
(i) Timing Analysis.
(ii) Analysis and Optimization of real-time and fault-tolerant applications implemented on distributed Platforms and MPSoC.
(iii) Analysis and Optimization of energy efficient, time constrained embedded systems.

Team leader: Prof. Lothar Thiele – TIK, ETH Zürich, ETHZ (Switzerland)
Developing a calculus to describe the performance of communication-centric systems, unifying the models for computation, combining tools for component-based performance analysis of MPSoC. Our role in this activity will be on component-based analytic methods to analyze the performance properties and memory requirements of distributed embedded systems.

Team leader: Prof. Rolf Ernst – TU Braunschweig, TUBS (Germany)
TU Braunschiweg contributes formal performance analysis methods for MpSoCs, with a focus on the timing implications of inter-task synchronization.

Team Leader: Prof. Maja D'Hondt – Interuniversity Microelectronics Centre, IMEC vzw. (Belgium) This team will introduce novel design-time and run-time resource management optimizations for MPSoC platforms.

Team leader: Prof. Luca Benini – University of Bologna, UNIBO (Italy)
(i) Development of power modeling and estimation framework for systems-on-chip.
(ii) Development of optimal allocation and scheduling techniques for energy-efficient mapping of multi-task applications onto multi-processor systems-on-chips.
(iii) Development of energy-scavenging techniques for ultra-low power sensor network platforms.

Professor, Prof. Axel Jantsch and Hannu Tenhunen, Royal Institute of Technolgy, KTH (Sweden)
The contribution form KTH focuses on various design aspects, architectures, run-time reconfigurability and adaptivity.

Team Leader : Dr. Raphaël David – CEA LIST (France)
(i) Development of exploration framework for multi- and many-core architectures
(ii) Development of advance strategies for the deployment and the management of multi-task applications onto multi- and many-core devices
(iii) Design of multi-core architectures for dynamic multi-task applications

Team Leader: Prof. Giovanni De Micheli – Swiss Institute of Technology, EPFL (Switzerland)
(i) Development of process-induced skew variability for clock distribution networks in 3-D ICS
(ii) Development of analytic thermal models for vertically integrated systems
(iii) Design automation tools for MPSoC and NoC
(iv) Optimization techniques for thermal management of MPSoC and NoC

-- Changes wrt Y2 deliverable --

IMEC vzw changed the Team leader. The new Team Leader is now Prof. Maja D'Hondt.

Prof. Giovanni De Micheli has been added as Team leader for EPFL.

1.2 Affiliated participants and their role within the Activity

Dr. Daniel Karlsson, Volvo Technology Corporation
Architecture and Design of Automotive Embedded Systems

Dr. Kai Richter – Symtavision (Germany)
Symtavision contributes industrial methodologies.

Dr. Arne Hamann – Robert Bosch GmbH (Germany)
Contributes on important embedded systems related research problems in the automotive industry.

Prof. Dimitrios Soudris – Democritus Uni. of Thrace, DUTH (Greece)
This team will introduce novel dynamic data type and data allocation optimizations for MPSoC platforms.

Prof. David Atienza – Uni. Complutense de Madrid, UCM (Spain)
This team will introduce novel run-time memory management optimizations for MPSoC platforms.

Prof. Per Gunnar Kjeldsberg - Norges Teknisk-Naturvitenskapelige Uni., NTNU (Norway)
This team will introduce novel task migration methodologies for MPSoC platforms utilizing hardware accelerators.

Bjørn Sand Jensen – Bang & Olufsen ICEpower (Denmark)
Areas of his team's expertise: chip design for audio signal processing

CTO Rune Domsteen – Prevas (Denmark)
Areas of his team's expertise: platform design for embedded systems

Dr. Patrick Schaumont – Virginia Tech (USA)
Design methods and architectures for secure embedded systems. Hardware/software codesign tool.

-- Changes wrt Y2 deliverable --

No changes with respect to Year 2.

1.3 Starting Date, and Expected Ending Date

Starting date: January 2008.

Ending date: Modeling and analysis is a long term effort and is expected to continue after the end of the project due to the lasting integration achieved by the NoE.

-- Changes wrt Y2 deliverable --

No changes with respect to Year 2.

1.4 Policy Objective

With growing maturity of scalable performance analysis algorithms and tools, new aspects such as the platform robustness can be included in analysis. Robustness to changes is particularly important for systems on chip since the cost of a redesign is high. At the same time robustness to faults is becoming a concern with shrinking feature sizes. In most practical cases, power consumption must be considered. There is currently no team in Europe that addresses all aspects. So integration of methods and tools will be needed to be able to (1) define meaningful robustness metrics that reflect design tradeoffs (2) assess the robustness of a design based on such metrics. This integration will extend the world leading position of Europe in the field of scalable formal performance analysis to hardware platform and MPSoC design.

-- Changes wrt Y2 deliverable --

No changes with respect to Year 2.

1.5 Background

The activity will be based on the complementary expertise of the participating partners in terms of Hardware Platform and MPSoC Analysis. In particular, the following areas are covered: Power modeling and analysis, power robustness assessment (University Bologna), platform performance modeling (University Braunschweig), analytical methods for reliability, performance and adaptability analysis of execution platforms (University of Denmark), reliability modeling, analysis and optimization (University Linköping), interfaces that communicate at run-time, aspects that are relevant for the efficiency of the run-time mapping components (IMEC, Belgium), simulation techniques and tools for NoC performance estimation and validation, interconnect and communication centric performance estimation techniques (KTH Sweden).

In addition, there have been already joint work and publications by some of the members of this activity, which will be used as a valuable starting point.

In more details, the above mentioned group has been working intensively on Power Modeling for SoC Platforms. In particular, they developed a virtual platform for power modeling of complex multi-core systems on chip. This platform can facilitate further integration among partners and associates, thanks to is flexibility and generality. In terms of “scheduling based energy optimization for energy-scavenging wireless sensor networks”, a novel scheduling strategy (called lazy scheduling) that is well suited to energy-harvesting systems operating under real-time constraints has been developed by ETHZ and University Bologna. It is the first result of this kind in this quickly growing research area and received a lot of attention in the scientific community.

At ETHZ, an open tool set is available that allows the performance analysis of distributed embedded systems and MPSoC. It is based on the concept of Modular Performance Analysis (MPA). In addition, there are first results available that connect this system to the MPARM simulation framework from University Bologna and the Symta/S analysis system from University Braunschweig/Symtavision.

-- Changes wrt Y2 deliverable --

No changes with respect to Year 2.

1.6 Technical Description: Joint Research

The major focus of the activity on Platform Analysis is to establish a set of models and analysis methods that (a) scales to massively parallel and heterogeneous multiprocessor architectures, (b) is applicable to distributed embedded systems as well, (b) allows for the analysis of global predictability and efficiency system properties and (c) takes the available hardware resources and the corresponding sharing strategies into account. Promising approaches are based on composable frameworks and treating resources as first class citizens in the analysis. Both, simulation-based and analytic methods will be combined. In addition, methods that focus on worst-case/best-case results as well as those based on stochastic models will be combined.

As a central ingredient of any analysis model, synchronization & communication abstractions are required for successfully deploying MPSoC hardware in embedded application domains. Efficiency is inherently related to both power and performance; hence it is an energy metric. In embedded systems, abstractions are acceptable only if they do not compromise efficiency. It also extremely important to take into account variabilities of both hardware fabrics and application workloads, which are deemed to rapidly increase. In particular, the above abstractions need to be embedded into a framework that allows to analyze the performance properties and memory requirements of distributed systems. In particular, we will focus on methods that satisfy composability properties and to lift the component-based methods as known from software design to interfaces that talk about resource interaction. In addition, we are interested in adding run-time adaptivity to systems while using efficient run-time estimation methods combined with distributed finite horizon control methods. Again, the focus is on predictability AND efficiency. Here, we will use the expertise that is available at ETH Zurich (Lothar Thiele) and University Bologna (Luca Benini), Hannu Tenhunen (KTH), Stylianos Mamagkakis (IMEC), Jan Madsen (DTU), TU Braunschweig (Rolf Ernst) are involved in this activity.

Another major challenge is to provide analysis tools and techniques to support the transitions between different abstraction levels in the design flow. Constraints should be communicated at design-time from one step to the next, taking into account the global effect that they will introduce in the system. Also, in order to ensure adaptivity of the system an interface should communicate at run-time the changes in the resource requests and the changes in the actual resource availability.

-- Changes wrt Y2 deliverable --

No changes with respect to Year 2.

1.7 Work achieved in Year 1 *(Jan-Dec 2008)*

A number of problems were tackled in Year 1, through several cooperations involving two or more partners. In the following paragraphs, we briefly summarize the problems tackled and the partners involved.

The Linköping group has addressed two major issues: *Timing analysis of distributed real-time systems*. In this context, the emphasis was on heterogeneous systems using various task scheduling policies and heterogeneous communication protocols with static and dynamic phases. Both formal and simulation based approaches were developed. *Analysis of fault tolerant real time systems considering various reliability requirements and fault tolerance mechanisms has been done*. In particular, the issue of transient faults has been considered. The basic effort was toward development of adequate scheduling algorithms. This work has been performed in cooperation with the DTU group.

As a central ingredient of any analysis model, synchronization & communication abstractions are required for successfully deploying MPSoC hardware in embedded application domains. Efficiency is inherently related to both power and performance; hence it is an energy metric. In embedded systems, abstractions are acceptable only if they do not compromise efficiency. It is also extremely important to take into account variability of both hardware fabrics and application workloads, which are deemed to rapidly increase. In particular, the above abstractions need to be embedded into a framework that allows to analyze the performance properties and memory requirements of distributed systems. In particular, ETHZ had its focus on methods that satisfy composability properties and to lift the component-based methods as known from software design to interfaces that talk about resource interaction. In the past year, ETHZ followed the following research directions: Together with University Bologna, ETHZ worked on combining the MPARM simulation framework with the performance analysis framework MPA (modular performance analysis). In particular, we attempted to set up the simulation environment in a way, that it follows the semantics of the analytic models and a comparison is possible. Together with University Braunschweig, ETHZ worked on the combination of the Symta/S symbolic analysis tool with the MPA (modular performance analysis framework). In particular, a tool coupling has been established that allows a seamless integration of both analysis methodologies. In addition, both research groups worked on a proper modeling of hierarchical event stream concepts.

To provide a formal performance analysis that captures the timing implications of multiprocessor systems on chip, the applicability of concepts from the analysis distributed systems is limited. A major difference lies in the use of common resources, either physically, such as a shared coprocessor or memory, or logically, such as a semaphore or a shared data structure in the memory. In ARTIST2 already first steps were taken towards addressing implications of a shared memory, which need to be extended in order to achieve the goals of this activity. We have taken steps towards generalization of the concepts from shared memory modeling to cover arbitrary shared resources. The approach chosen promises a higher accuracy than traditional approaches, due to more sophisticated modeling of shared resource load, and a better composability of designs, as the analysis of the shared resource delay is decoupled from the response time analysis of the requesting tasks.

The main problem tackled by IMEC and its affiliated partners (ie, DUTH, NTNU and UCM) was the definition of a specific software metadata format, which can be linked to each embedded software application or downloadable software service. This software metadata information can be used to configure and self-adapt the run time resource management software for dynamic data transfer and storage on MPSoC platforms. Additionally, IMEC has developed profiling and analysis tools that extract and post-process these software metadata, in order to be used for both memory hierarcy assignment (ie, MH tool) and source code parallelization tools (ie, MPA tool).

University of Bologna has addressed, in cooperation with DTU, the issue of prolonging the system life-time. Even though systems that harvest energy from the environment are adopted, such an environmental energy is not distributed uniformly and there is a lot of parameters that influence the efficiency and the schedulability of tasks. In particular we tackled the problem of routing messages in a sensor network with energy awareness and real-time responsiveness together with scheduling policies, which guarantee to execute tasks under unpredictable profile of the harvested energy.

-- No changes wrt Y2 deliverable --

This section was already presented in the Y2 deliverable, in section 1.7.

1.8 Work achieved in Year 2 *(Jan-Dec 2009)*

Many of the problems tackled in Year 2 are continuations of the problems and cooperations initiated in Year 1.

Linköping and DTU have continued their work regarding analysis and optimization of distributed embedded real-time systems, with application in automotive electronics. The issue of average response time (as a complement to the previous work, regarding worst case response times) of FlexRay-based distributed systems has been addressed.

At Linköping, in cooperation with Lund, analysis and optimisation techniques for control–scheduling co-design have been developed. It integrates a controller design with both static and priority-based scheduling of the tasks and messages, and in which the overall control performance is optimized. The technique has also been extended to cover the case of multimode control systems.

As a main result during the second year, the groups at Linköping and DTU have developed an analysis approach that determines the system failure probability, based on the number of re-executions introduced in software and the process failure probability of the computation node with a given hardening level. Based on this analysis, an optimisation technique has been proposed in which hardware and software fault tolerance techniques are combined.

TU Braunschweig and GM Labs have collaborated in the COMBEST project on the definition of methods and tools for the timing analysis of automotive systems based on a mix of complex communication protocols/software scheduling techniques. Within this area, ETHERNET (and its different implementations, e.g., AFDX and TT-ETHERNET) was of special interest, because increasing bandwidth, low latency/jitter, and time determinism requirements make an Ethernet based solution for in-vehicle networking attractive. In this scope, the partners examined the application of methods and tools to possible future use cases from the automotive domain. To examine the accuracy of the conservative performance bounds obtained by the formal worst case analysis, a comparison between results of formal analysis and results obtained by simulation was performed. A joint publication containing the results of the comparison has been submitted for publication.

UNIBO and Linköping have tested Basic Scheduling Alternative (BSA) in a Multi-Processor System-on-Chip shared bus. A RTL-level cycle accurate TDMA bus arbiter model has been developed and plugged in MPARM simulator. Many exhaustive sets of experiments have been done in order to validate the BSA framework and during each step of the exploration, a hardware synthesis has been performed to keep under control the underlying logic complexity.

DTU has continued the work on analytical performance models. In particular, the formalisation of the ARTS model has been extended to capture more aspects of both the application and the platform. Another important issue to be addressed is the scalability of the model such that larger and more complex systems can be modelled.

During the annual Meeting of the ArtistDesign Cluster on Hardware Platforms and MpSoC in Braunschweig, the partners of this activity have identified common goals and similarities in their approaches to address shared resources in multiprocessor systems. Shared resources generally complicate the timing verification, because they introduce timing inter-dependencies between the tasks on different cores. To address this challenge, Linköping University, ETH Zürich, TU Braunschweig are working on formal approaches to increase the predictability of such systems: Linköping Universitys focus is finding static bus schedules that minimize the overhead from a worst-case perspective. At the ETH Zürich, work is being done on finding the worst-case timing interference within assigned time-slots. The TU Braunschweig is working on an approach that allows to bound the interference in the absence of static schedules.

In the collaboration between ETHZ and University of Bologna, further investigation of energy harvesting systems was performed. In particular, this concerned both hardware and software asptects of sensor nodes which are powered solar energy. The main goal was to merge the different approaches and illuminate several application scenarios that are of practical relevance.

Concerning the integration of SymTA/S and MPA as well as unifying approaches for hierarchical scheduling, hierarchical event streams have been in the focus of research in the second year. The major issue here is to also allow an extraction of single event streams from previously merged and transformed event streams. The transformation is hereby due to the fact that incoming streams can be combined via OR-operation and may pass different system components, which may buffer these streams due to scheduling policies.

CEA LIST has been involved in a new cooperation with university of Bologna for the definition of a Software Runtime Architecture for the management of many-cores components. CEA LIST has proposed with University of Bologna a framework supporting the development of various services, ranging from dynamic application deployment, task scheduling, resources allocation to fault and power management.

IMEC focused its research closer to the hardware platform and more specifically regarding variability and reliability issues arising from the usage of sub 45nm technology nodes. To this end Variability Aware Modeling focused on analysis of SRAM designs and the Knobs and Monitors framework extensions focused on RTL2RTL transformations for latency monitors.

Finally, KTH and DTU have started an collaboration within the SYSMODEL project on the development of a multi-MoC modeling framework for heterogeneous systems integrated with performance analysis and design space exploration tools.

-- No changes wrt Y2 deliverable --

This section was already presented in the Y2 deliverable, in sections 1.8 and 3.1.

1.9 Problems Tackled in Year 3 *(Jan-Dec 2010)*

University of Bologna (**UNIBO**) tackled the scheduling problem on multicore systems An effective multicore computing is not only a technology issue, as we need to be able to make efficient usage of the large amount of computational resources that can be integrated on a chip. This has a cross-cutting impact on architecture design, resource allocation strategies, programming models.The objective of our work is the development of a predictable and efficient non-preemptive scheduling of multi-task applications with inter-task dependencies.

Moreover **UNIBO** addressed the thermal and reliability managment for multicore platforms together with **Intel Labs**. (Braunschweig). The goal is to mitigate the onchip temperature hotspots, which can cause non-uniform ageing, by online tuning of system performance and adopting closed-loop controls.

ETHZ has been investigating various component-based approaches to the analysis of MPSoC. In particular, new interface-based approaches have been proposed that allow for an end-to-end analysis based on a novel interface algebra. The work is described in [SCT10].

ETHZ and **UNIBO** have been pushing forward the adaptive strategies for power management in case of environmentally powered systems. This joint work has been published in [MTBB10].

Furthermore, **ETHZ** has been investigating various methods to analyse the performance of multiprocessor systems under power and temperature constraints, e.g. involving leakage power. Finally, the component-based approach has been extended towards combining various

models of computation, i.e. state-based component descriptions and analytic functional models, see [LPT10]. There has been substantial progress in understanding the influence of memory buses on the performance of multicore systems, see [SPCTC10].

The collaboration between **TU Braunschweig** and GM Labs in the COMBEST project started in 2009 was focused on the application of methods and tools developed in COMBEST on possible future test use cases from the automotive domain. The results of this collaboration have been published at SAE 2010 [RE10]. As a follow up to the successful collaboration in 2009, in 2010, two separate R&D projects between iTUBS and GM Labs have been initiated to further exploit the COMBEST research results. The focus of these projects was on the development of methods and tools for the analysis (i) of Ethernet based architectures and (ii) of multi-core platforms. Both projects, which are currently still ongoing, have a strong focus on applying formal scheduling theory to realistic foreseeable automotive architectures. In particular, the collaboration with GM Labs shows that the techniques developed in COMBEST (e.g. formal performance analysis of Ethernet based communication networks) help to obtain formal analysis results for realistic use cases of GM and give much tighter analysis results than the formal techniques previously available. This generally increases the user acceptance of formal analysis methods and it also increases the user value of the tool SymTA/S.

TU Braunschweig has continued the work on the performance analysis for multiprocessor systems with shared resources. Timing implications of using shared resources in multi-core systems have been investigated for different setups and, key components of a safe inter-task and inter-core synchronization algorithm for shared resources have been identified. Also, in a joint meeting **TU Braunschweig** and **ETHZ** have discussed open issues and identified possible improvements on the modelling and analysis approaches of multi-core systems with shared resources. Results have been presented jointly by **TU Braunschweig** and **ETHZ** at the annual ArtistDesign Cluster Meeting in Leuven. During this meeting, further ideas regarding approaches addressing shared resources in multiprocessor systems have been exchanged with partners of this activity, e.g. **Linköping University.**

Concerning contract based architecture dimensioning, **KTH** has formulated flow based traffic shaping optimization problem for minimizing buffers in the communication network, while meeting the throughput and delay requirements for each flow. In a case study we observed 62.8% reduction of total buffers, 84.3% reduction of total latency, and 94.4% reduction on the sum of variances of buffers. Likewise, the experimental results demonstrate similar improvements in the case of synthetic traffic patterns. The optimization algorithm has low run-time complexity, enabling quick exploration of large design spaces. From those experiment it can be concluded that flow regulation can be a valuable instrument for buffer optimization in NoC designs.

The memory organization and the management of the memory space is a critical part of every NoC based system. **KTH** has developed a Data Management Engine (DME), that is a block of programmable hardware and part of every processing element. It off-loads the processing element (CPU, DSP, etc.) by managing the memory space, memory access and the communication over the on-chip network. The DME's main functions are virtual address translation, private and shared memory management, cache coherence protocol, support for memory consistency models, synchronization and protection mechanisms for shared memory communication. The DME is fully programmable and configurable thus allowing for customized support for high level data management functions such as dynamic memory allocation and abstract data types.

A system modelling framework for modelling, analysis and refinement of heterogeneous MPSoCs is being developed by **KTH** in cooperation with **DTU** and Tampere University. The framework rests on the formal basis of ForSyDe and is extended for modelling of both the functionality and platform architecture. Also, a SystemC implementation has been defined and is under development. To allow the modeling of abstract heterogeneous system models we

have speficied the System Functionality Framework (SFF), which defines the modeling of functionality at a high-level of abstraction. The SFF comprises several models of computation, which are formally defined using the ForSyDe framework. SystemC has been selected as the main modeling language for the SFF. The Platform Architecture Framework (PAF) is a library of platform components modeled at different levels of abstraction. PAF includes processor cores, memory models, interconnection structures, and instruction set simulators and other components. The system level models serve as behavioral references (or executable specifications) for the lower level components. Therefore, the system level models alleviate the design validation and design space exploration.

The **CEA LIST** and **UNIBO** focused during year 2010 on the exploration and the development of dynamic management strategies for the deployment and the execution of multi-tasks application on many-cores architectures. Special attention has been paid on load-balancing and power management. Also, smart deployment of application, taking care of potential hardware defauts in the device, has been explored.

Temperature aware system level design methods rely on the availability of temperature modeling and analysis tools. The **Linköping** group has developed new, fast, and accurate temperature analysis techniques that are sufficiently fast to be and used inside a temperature aware system level optimization loop.

Linkoping and **Lund** has worked together on QoS optimization of real-time control applications. Time-triggered periodic control implementations are over provisioned for many execution scenarios in which the states of the controlled plants are close to equilibrium. To address this inefficient use of computation resources, researchers have proposed self-triggered control approaches in which the control task computes its execution deadline at runtime based on the state and dynamical properties of the controlled plant. The potential advantages of this control approach cannot, however, be achieved without adequate online resource-management policies. The groups at Linköping and Lund have addressed the issue of performance modelling and resource scheduling for multiple self-triggered control tasks that execute on a uniprocessor platform.

DTU has continued the work on analytical performance models. In particular, the formalisation of the ARTS model has been extended to capture more aspects of both the application and the platform. Another important issue to be addressed is the scalability of the model such that larger and more complex systems can be modelled.

An important part of the work implemented by **EPFL** is concerned with the synthesis of application specific Network-on-Chip (NoC) topologies for 3D integrated systems [SMBM10]. The work present a complete algorithm with two underlining approaches that is capable to explore the tradeoffs between topology performance (power consumption, network delay) and the number of required vertical links that use Through Silicon Vias (TSVs). The synthesis algorithm is extended with a custom floorplanning routine that preserves the relative positions of the core blocks. The work also presents a comparison between application specific NoC for 2D and 3D Integrated circuits (ICs) and presents the benefits of the topologies for the 3D ICs.

A further study makes a comparison between NoC topologies for 2D and 3D IC in the presence of Voltage and Frequency Islands (VFI) [SMBM10b]. The study shows the advantages of 3D topologies for different VFI assignments of the System –on-Chip (SoC) cores. The work also shows the need for synthesis tools for the exploration and design of custom topologies for 3D ICs.

Deadlock freedom is a necessity for NoC. Traditional deadlock avoidance techniques that restrict the synthesis tools can lead to the impossibility of the synthesis algorithm to find paths in the presence of additional constraints (e.g. restricted number of vertical links that require TSVs) [SMBM10c]. We presented an algorithm to remove possible deadlocks after synthesis. The algorithm can be applied to any topology and removes the deadlocks by minimally adding

parallel physical links or virtual channels. This work was done in cooperation with **UNIBO**. The goal of the collaboration is to provide tools for application specific NoCs for realistic systems that use VFIs (since most SoCs do today) and for 3D IC (which is a promising technology for future SoCs).

Thermal management management policies for MPSoCs have also been another focus within **EPFL**. This project was a partly joint effort with **ETHZ**. Three-Dimensional integrated circuits and systems are expected to be present in electronic products in the short term. Specifically 3-D multi-processor systems-on-chip (MPSoCs), realized by stacking silicon CMOS chips and interconnecting them by means of through silicon vias has been considered. Because of the high volumetric density of devices and interconnect, thermal issues pose critical challenges, such as hot-spot avoidance and thermal gradient reduction. Thermal management is achieved by a combination of active control of on-chip switching rates as well as active interlayer cooling with pressurized fluids.

-- The above is new material, not present in the Y2 deliverable --

2. Summary of Activity Progress in Year 3 *(Jan-Dec 2010)*

2.1 Technical Achievements

Power, Thermal and Reliability Aware Resource Managment for Multicore Systems (University of Bologna, Intel Labs)

The use of high-end multicore processors today can incur high power density with significant variability in spatial and temporal usage of resources by workloads. This situation leads to power and temperature hotspots, which in turn may lead to non-uniform ageing and accelerated chip failure. These drawbacks can be mitigated by online tuning of system performance and adopting closed-loop thermal and reliability management policies. The development and evaluation of these policies cannot be performed solely on real hardware - due to observability and flexibility limitations or just by relying on trace-driven simulation, due to dependencies present among power, thermal effects, reliability and performance.

University of Bologna developed a complete virtual platform to explore, simulate and evaluate power, temperature and reliability management control strategies for multicores platform. The accuracy and effectiveness of our solution are ensured by integrating a established system simulator with models for power consumption, temperature distribution and aging. The models are based on characterization on real hardware. Control strategies exploration and design are carried out in the MATLAB/Simulink framework allowing the use of control theory tools. Fast prototyping is achieved by developing a suitable interface between functional simulator and MATLAB/Simulink, enabling co-simulation of hardware platforms and controllers. This research activity has been pursuit with in collaboration with Matthias Gries of the ***Intel Lab***, German Microprocessor Lab (Braunschweig, Germany).

Scheduling for conditional task graphs in multi-processor systems-on-chip (UNIBO)

Multicore platforms have become widespread in embedded computing. This trend is propelled by ever-growing computational demands of applications and by the increasing number of transistors-per-unit-area, under continuously tightening energy budgets dictated by technology and cost constraints. Effective multicore computing is however not only a technology issue, as we need to be able to make efficient usage of the large amount of computational resources that can be integrated on a chip. This has a cross-cutting impact on architecture design, resource allocation strategies and programming models.

The objective of **University of Bologna** (**UNIBO**) is predictable and efficient non-preemptive scheduling of multi-task applications with inter-task dependencies. UNIBO focuses on non-preemptive scheduling because it is widely utilized to schedule tasks on clustered domain-specific data-processors under the supervision of a general-purpose CPU. The target platform is described as a set of resources having finite capacity. An application is a set of dependent tasks that have to be periodically executed with a max-period constraint (i.e. deadline equal to period), modeled as a task graph where task execution is characterized by known min-max intervals and inter-task dependencies are associated with a known amount of data communication. A preliminary heuristic allocation step is assumed to assign tasks to processor clusters targeting workload balancing and minimization of inter-cluster communication; task migration among clusters is not allowed.

The achievement is providing off-line a partitioned scheduling, guaranteeing to meet hard real-time constraints. **UNIBO** developed a robust scheduling algorithm that proactively inserts additional inter-task dependencies only when required to meet the deadline for any possible combination of task execution times within the specified intervals. The used approach is to

avoid idle time insertion, hence it does not artificially lower platform utilization, it does not need timers and related interrupts, nor a global timing reference. The algorithm is complete, i.e. it will return a feasible graph augmentation if one exists.

UNIBO also proposes an iterative version of the algorithm for computing the tightest deadline that can be met in a robust way. Even though worst-case runtime is exponential, the algorithms have affordable run times (i.e seconds to minutes) for task graphs with tens of tasks.

Component-based analysis of complex MPSoC Systems (ETHZ, UBS)

There has been substantial progress in understanding the component-based analysis of complex MPSoC systems. The results are described in publications [SCT10], [LDT10], [PRTLT10]. Interface-based design relies on the idea that different components of a system may be developed independently and a system designer can connect them together only if their interfaces match, without knowing the details of their internals.

In [SCT10] we propose an interface algebra for analyzing networks of embedded systems components. The goal is to be able to compute worst-case traversal times and verify their compliance to provide deadlines in such component networks in an incremental manner, i.e., as and when new components are added or removed from the network.

[LDT10] presents a compositional and hybrid approach for the performance analysis of distributed real-time systems. The developed methodology abstracts system components by either flow-oriented and purely analytic descriptions or by state-based models in the form of timed automata. The interaction among the heterogeneous components is modeled by streams of discrete events. In total this yields a hybrid framework for the compositional analysis of embedded systems. It supplements contemporary techniques for the following reasons: (a) state space explosion as intrinsic to formal verification is limited to the level of isolated components; (b) computed performance metrics such as buffer sizes, delays and utilization rates are not overly pessimistic, because coarse-grained analytic models are used only for components that conform to the stateless model of computation. For demonstrating the usefulness of the presented ideas, a corresponding tool-chain has been implemented. It is used to investigate the performance of a two-staged computing system, where one stage exhibits state-dependent behavior that is only coarsely coverable by a purely analytic and stateless component abstraction.

[PRTLR10] extends the methodology of analytic real-time analysis of distributed embedded systems towards merging and extracting sub-streams based on event type information. For example, one may first merge a set of given event streams, then process them jointly and finally decompose them into separate streams again. In other words, data streams can be hierarchically composed into higher level event streams and decomposed later on again. The proposed technique is strictly compositional, hence highly suited for being embedded into well known performance evaluation frameworks such as Symta/S and MPA (Modular Performance Analysis). It is based on a novel characterization of structured event streams which we denote as Event Count Curves. They characterize the structure of event streams in which the individual events belong to a finite number of classes. This new concept avoids the explicit maintenance of streamindividual information when routing a composed stream through a network of system components. Nevertheless it allows an arbitrary composition and decomposition of sub-streams at any stage of the distributed event processing.

Power Management in Environmentally Powered Systems (ETHZ, UNIBO)

The joint work described in [MTBB10] describes a novel approach to the analysis of environmentally powered systems. Based on a prediction of the future available energy, we adapt parameters of the application in order to maximize the utility in a long-term perspective.

The paper presents a formal model of the corresponding optimization problem including constraints concerning buffer sizes, timing, and rates. Instead of solving the optimization problem online which may be prohibitively complex in terms of running time and energy consumption, we apply multiparametric programming to precompute the application parameters offline for different environmental conditions and system states.

MPSoC System Bus Analysis (ETHZ, UNIBO)

Multi-processor and multi-core systems are becoming increasingly important in time critical systems. Shared resources, such as shared memory or communication buses are used to share data and read sensors. In [SPCTC10], we consider real-time tasks constituted by superblocks, which can be executed sequentially or by a time triggered static schedule.Three models to access shared resources are explored: (1) the dedicated access model, in which accesses happen only in dedicated phases, (2) the general access model, in which accesses could happen at anytime, and (3) the hybrid access model, combining the dedicated and general access model. For resource access based on a Time Division Multiple Access (TDMA) protocol, we analyze the worst-case completion time for a superblock, derive worst-case response times for tasks and obtain the relation of schedulability between different models. We conclude with proposing the dedicated sequential model as the model of choice for time critical resource sharing multi-processor/multi-core systems.

Modelling and Analysis of Multiprocessor Systems with Shared Resources (TU Braunschweig)

In year 3, TU Braunschweig has continued the work on developing analysis methods for Ethernet based communication networks. These methods have been used to investigate timing characteristics of Ethernet-based communication network solutions expected to be implemented in future automotive systems. The work on comparing the results obtained with the developed formal analysis methods with simulation results is ongoing in collaboration with GM Labs. Results have not yet been published.

TU Braunschweig has further worked on modelling and analysis methods for multiprocessor systems with shared resources. A method that captures more accurately the interference between different cores that share common resources [SNE10] has been applied to different multi-core setups to investigate the advantages on the obtained analysis results. Further the modelling and analysis framework has been extended to investigate the impact of different design options of the shared resource arbitration mechanisms on the systems timing. For this, solutions proposed by academia, e.g the MPCP – Multiprocessor Priority Ceiling Protocol, and by industry, i.e. arbitration strategy proposed in the last released AUTOSAR specification have been considered. Results have been published in the paper [NSE10] that has been later selected for a journal [NSE+10]. Further work was performed to investigate the applicability of the analysis methods on particular problems in the automotive domain. The work is ongoing, solutions and results have been formulated but not yet published.

Contract based architecture dimensioning (KTH)

Defining and constraining traffic in the on-chip network can be a means to achieve predictable performance at low cost. Based on traffic contracts between IPs and the communication infrastructure the network can be optimized to meet all performance constraints at minimum cost. For network-on-chip (NoC) designs, optimizing buffers is an essential task since buffers are a major source of cost and power consumption. KTH has developed a flow regulation technique and has defined a regulation spectrum as a mean for system-on-chip architects to

control delay and backlog bounds. The regulation is performed per flow for its peak rate and burstiness. However, many flows may have conflicting regulation requirements due to interferences with each other. Based on the regulation spectrum, this flow optimizes the regulation parameters aiming for buffer optimization. Three timing-constrained buffer optimization problems are formulated, namely, buffer size minimization, buffer variance minimization, and multiobjective optimization, which has both buffer size and variance as minimization objectives. Minimizing buffer variance is also important because it affects the modularity of routers and network interfaces. A case study exhibits 62.8% reduction of total buffers, 84.3% reduction of total latency, and 94.4% reduction on the sum of variances of buffers. Likewise, the experimental results demonstrate similar improvements in the case of synthetic traffic patterns. The optimization algorithm has low run-time complexity, enabling quick exploration of large design spaces. Thus, flow regulation can be a valuable instrument for buffer optimization in NoC designs.

Integration of the communication architecture with the memory architecture (KTH)

KTH has developed a Data Management Engine (DME) that manages all communication and a distributed, shared memory space in an MPSoC. The DME if fully programmable and offers the following functionalities:

- Virtual memory space
- Cache coherency
- Memory consistency
- Message passing communication
- Shared memory based communication
- Dynamic memory allocation

The DME is a fully programmable controller but optimized for memory management with specific HW support for a number of specialized functions. It requires approximately 50 kGates plus control store for implementation and can run at 500MHz in a 90nm process. A configuration tool and an assembler support the customization of the DME. Experiments with applications, that inherently allow for parallelization, have shown that a DME based implementation with effective distributed shared memory managent, can perform 10x better than a traditional implementation based on a central memory.

Modeling and analysis of heterogeneous systems (KTH, DTU)

As part of the ARTEMIS project SYSMODEL KTH, DTU and Tampere University are developing a complete system-level modelling framework for analysis and refinement of heterogeneous MPSoCs. During 2010 a System Functionality Framework (SFF) and a Platform Architecture Framework (PAF) have been specified. The implementation in SystemC has started and is expected to complete at the end of 2011. Furthermore we have defined the Design Space Exploration (DSE) methodology for mapping SFF models and PAF structures into an implementation. The design framework allows to model heterogeneous systems at a high abstraction level. The SFF supports several models of computations, which allows to model systems consisting of software, and digital and analog hardware. The SFF has a precise semantics facilitating the application of formal methods. SystemC, an industrial design language, is used as the main modelling language. However, it is also possible to include synthesizable VHDL models and algorithms specified in C or MATLAB into the SFF. The PAF is mainly a library of platform components, which are modeled at different levels of abstraction.

The design framework allows to co-simulate SFF system-level processes together with platform components, which are encapsulated in a SystemC wrapper. This allows a stepwise refinement by replacement, where processes of the abstract SFF model are replaced by platform components. For the simulation of the refined model the same SystemC test-bench that has been used for the original SFF model can be used. Refinement by replacement is accompanied by design space exploration methods, which exploit the formal nature of the SFF model and the performance figures that are provided by the platform components.

Runtime layer design for many-cores architectures (CEA, UNIBO)

CEA LIST has implemented a software environment that allows firmware developers to master the inherent complexity of parallel programming for manycore architectures. It does not assume a unique programming model; instead several are made available to better suit programmers' needs. This is done possible by the definition of several basic execution primitives (tasks and threads). Threads fit well with coarse grain parallelization processes, whereas lightweight threads or tasks are perfectly suited for fine-grain parallelism like that of work-items of OpenCL for example. Both taks and threads are dynamically managed to load balance the fabric at runtime.

The runtime SW provides a low-level virtualization of the whole fabric that allows a complete decoupling of the logical allocation process (selecting computation, storage and communication resources) from the physical assignment. This low-level virtualization layer allows to deploy applications independently from the underlying available hardware and thus potentially improve manufacturing yield. Both error and power management services are implemented using a two-step approach, leveraging a closed-loop control of the resources and a global optimization strategy. This hierarchical approach combines high reactivity to local events, and high efficiency of global optimization processes. On the power management side, it may allow for example to quickly switch into idle mode as soon as a processor is stalled (waiting for a new job to be done, a data to be sent, etc.), while a global optimization process deals with task allocation to maximize idling time and reduce the idle mode toggling frequency. On VC1-decoding applications CEA and UNIBO has shown improvement of 25% in energy efficiency.

To support a wide-diversity of utilization contexts and programming models, special attention has been given to the design of the SW runtime architecture to allow a seamless composition of the services to be supported. The runtime software has thus been split into several modules according to the role of each one, with strict module dependency rules. Each module is compiled into a shared library so as to enable code sharing between applications at load time, and make use of only code portions required by the application. On the basis of a generic kernel, with a set of features independent of the programming model, several applications can therefore be simultaneously loaded with different execution models and QoS.

System Level Temperature Modeling and Analysis (Linkoping)

Temperature aware system level design methods rely on the availability of temperature modelling and analysis tools. System level temperature modelling approaches are mostly based on the duality between heat transfer and electrical phenomena. However, temperature analysis time with available approaches, such as *Hotspot* or *ISAC* like are too long to be affordable inside a temperature aware system level optimization loop. There has been some work on establishing fast system level temperature analysis techniques. They also build on the duality between heat transfer and electrical phenomena and are based on very restrictive assumptions in order to simplify the model. Such typical assumptions are that (1) no cooling layer is present, (2) there is no interdependency between leakage current and temperature,

and (3) the whole application executes at a constant voltage. In order to support our temperature aware energy optimisation techniques, the Linköping group proposed a fast and accurate temperature analysis approach which eliminates all three limitations mentioned above and can be used inside a temperature aware system level optimization loop.

Modeling and analysis for QoS optimisation of real-time control applications (Linkoping, Lund)

Time-triggered periodic control implementations are over provisioned for many execution scenarios in which the states of the controlled plants are close to equilibrium. To address this inefficient use of computation resources, researchers have proposed self-triggered control approaches in which the control task computes its execution deadline at runtime based on the state and dynamical properties of the controlled plant. The potential advantages of this control approach cannot, however, be achieved without adequate online resource-management policies. The Linköping and Lund groups have addressed the issue of QoS modelling and scheduling for multiple self-triggered control tasks that execute on a uniprocessor platform, where the optimization objective is to find tradeoffs between the control performance and CPU usage of all control tasks. the results show that efficiency in terms of control performance and reduced CPU usage can be achieved with the heuristic proposed in this paper.

A self-triggered control task computes deadlines on its future executions, by using the sampled states and the dynamical properties of the controlled system, thus cancelling the need of specialized hardware components for event generation. The deadlines are computed based on stability requirements or other specifications of minimum control performance. Since the deadline of the next execution of a task is computed already at the end of the latest completed execution, a resource manager has, compared to event based control systems, a larger time window and more options for task scheduling and optimization of control performance and resource usage. In event-based control systems, a control event usually implies that the control task has immediate or very urgent need to execute, thus imposing very tight constraints on the resource manager and scheduler. The contribution of this work is a software-based middleware component for scheduling and optimization of control performance and CPU usage of multiple self-triggered control loops on a uniprocessor platform. The proposed technique is based on cost-function approximations and search strategies. Stability of the control system is guaranteed through a design-time verification and by construction of the scheduling heuristic.

Modeling and analysis of fault tolerant distributed embedded systems (DTU, Linkoping)

DTU has focused on mixed-criticality systems. The “Research Agenda for Mixed-Criticality Systems” [1] defines a mixed-criticality system as “an integrated suite of hardware, operating system and middleware services and application software that supports the execution of safety- critical, mission-critical, and non-critical software within a single, secure computing platform”. The current practice to mixed-criticality systems is to physically separate the different criticality functions in different hardware components, so they cannot influence each other. In avionics, the proposed integration solution is based on “Integrated Modular Avionics” (IMA), which allows the integration of mixed-criticality functions onto the same node as long as there is enough spatial and temporal partitioning. A similar problem is faced in many other industries. This is coupled with the trend towards using multi-core systems, where several processors can be integrated onto a single chip, decreasing the costs, power consumption, size, and increasing the performance through parallelization. At the European level, multi-cores are addressed through EU projects such as RECOMP (“Reduced certification cost for trusted multi-core platforms”), which has the goal to “define a European standard reference technology for mixed- criticality multi-core systems supported by the European tool vendors”.

In the context of mixed-criticality systems, DTU has done an initial investigation on the optimization of temporal partitions. The assumption is that the architecture provides both spatial and temporal partitioning, thus enforcing enough separation for the critical applications. With temporal partitioning, each application is allowed to run only within predefined time slots, allocated on each processor. The sequence of time slots for all the applications on a processor are grouped within a Major Frame, which is repeated periodically. We have assumed that the safety-critical applications (on all criticality levels) are scheduled using static-cyclic scheduling and the non-critical applications are scheduled using fixed-priority preemptive scheduling. Each application runs in a separate partition, and each partition is allocated several time slots on the processors where the application is mapped. We have proposed a Simulated Annealing-based optimization meta-heuristic to determine the sequence and size of the time slots within the Major Frame on each processor such that both the safety-critical and non-critical applications are schedulable. The results have been published in a paper submitted to AMICS 2011, The 1st International Workshop on Architectures and Applications for Mixed-Criticality Systems.

DTU and Linköping have collaborated on the impact of time partitioning on the analysis of tasks scheduled using a fixed-priority preemptive scheduling policy. Linköping has proposed a schedulability analysis for such tasks, which extends the schedulability analysis with static and dynamic offsets to consider the influence of statically-scheduled non-preemptive tasks on the worst-case response times of the fixed-priority preemptive tasks. Together with Linköping, DTU has extended this analysis to consider the influence of time-partitions on the schedulability of the fixed-priority preemptive tasks. They have introduced the notion of demand and availability, used to compute the length of the busy window, which is needed in order to compute the worst-case response time of a task (the worst-case response time is then compared to the deadline). DTU has extended the concepts of demand and availability to take into account the time-partitions, and has integrated this analysis in the optimization of the temporal partitions.

Modeling and Verification of Embedded Systems (DTU, AAU)

One of the major challenges in designing an embedded system is to find a mapping of the application onto the execution platform which effectively fulfills the non-functional requirements of the embedded system such as timing, memory usage, energy consumption, and other cost. A particular challenge is to model and analyze cross-layer dependencies, where the change of a property in one part of the system, e.g. scheduling policy, may impact the performance of another part of the system, e.g. deadline miss on another processor, and hence, the overall system performance. The ARTS simulation model developed by DTU, has been modeled using timed automata and implemented in UPPAAL from AAU.

In order to make the formal model available for easy adaptation of embedded systems designers, the UPPAAL based model has been embedded in a tool called MOVES. MOVES supports formal analysis of non-functional properties of an embedded system, covering the system layers of an application mapped on an execution platform, consisting of a heterogeneous multiprocessor architecture where each processor may run a real-time operating system, and where all processors are connected through a network. It supports the designer by allowing him/her to describe the application, the execution platform and the mapping in a straight forward manner. MOVES then translates the system into a UPPAAL model which is then used to model check the system against given properties. If the model checking fails, the given counterexample produced by UPPAAL, is translated by MOVES into a schedule indicating where the properties were violated. The designer can then use this information to understand why the system failed and to suggest improvements.

DTU has continued the work on analytical performance models. In particular, the formalisation of the ARTS model has been extended to capture more aspects of both the application and the platform. The MoVES framework for analysis of embedded systems, developed by DTU, has

been streamlined, theoretically and practically, through the completion of the phd-thesis by Aske Wiid Brekling. Furthermore, the MoVES modelling language has been put in a context with the hardware verification language GEZEL.

Formal verification of design properties of hardware architectures (DTU, Virginia Tech)

As the complexity of chips grows, the methodology to build chips has to evolve. Today, chips are largely synthesized from high-level architectural descriptions that hide low-level details.

The majority of hardware designs are done using the most common hardware description languages, such as VHDL or Verilog. Both languages support high-level architectural descriptions, but allow hardware designers to incorporate low-level details in order to optimize for a particular hardware technology and directly synthesize using a restricted subset of the languages. However, chips may also be synthesized from software based models in much the same way as compilers produce executable code. Examples of such languages are Esterel, Lustre and Signal.

DTU have developed a language for hardware models based on the Gezel hardware description language developed and maintained by Virginia Tech, USA. The language depends on reasonably few, simple and clean concepts, and it strikes a balance between software and hardware concerns that suits the needs for a modern top-down approach to hardware design.

DTU have given a semantics domain that can be used for hardware design languages like Gezel. They have shown how the semantics can be used in connection with verification by relating the semantic domain to timed-automata using the UPPAAL system. A few simple example circuits have been successfully model and verified, e.g. the Simplified Data Encryption Standard (SDES) Algorithm and different algorithmic implementations of the Greatest Common Divisor.

DTU has started to investigate how this platform level modeling can be linked with the system-level models. Initial studies between DTU and Virginia Tech towards using the Gezel model for education in Embedded Systems Design has been conducted. Patrick Schaumont from Virginia Tech has written a text book, which among others has been used at DTU. Finally, DTU organized a half-day tutorial on Gezel and its usage for embedded systems education at the NORCHIP 2010 conference in Tampere, Finland.

Analysis tools for embedded systems (DTU, Oldenburg)

The goal is to establish efficient methods for verification of resource constraints, where one activity will focus on the real-time logic durations calculus. The aim of DTU is to develop a prototype model-checking tool and experiment with verifying strong timing constraints.

A refinement of the model checking theory and first prototype model-checker has been completed. This, second prototype implementation uses more fine-grained approximations and more compact data-structures than the first approximations. Furthermore, it uses the Z3 SMT-solver as back-end verifier for Presburger arithmetic. The Z3-solver is being replaced by an "in-house" solver for Presburger arithmetic, where quantifier-elimination techniques special geared to the problem at hand is developed and used.

Synthesis Tool for 3-D NoCs (EFFL, UNIBO)

Since EPFL has joined the activity only last year, new and recent topics and, consequently achievements are described herein. The activities performed this year are related to these problems and resulting publications described in Sections 1.9 and 2.2, respectively. A primary achievement of this year in collaboration with UNIBO was a synthesis tool for NoC systems that exploit the third physical dimension based on the floorplan of the processing elements of

the NoC. Furthermore, the alleviate thermal problems that can emerge in these 3-D NoC systems DVFS schemes are an efficient means to limit the development of hot spots. To this end, a comparative analysis between 2- and 3-D NoC based SoCs has been performed where multiple voltage and frequency islands are employed.Again, this piece of work has been performed in collaboration with UNIBO

Thermal Models and Thermal Management Schemes (EPFL, ETHZ)

In order to address the challenges pertinent to the design of 3-D NoC (which can be synthesized), specifically, thermal issues compact thermal models for the vertical interconnects have been developed. The proposed models as compared to the conventional approaches consider additional paths of heat transfer that considerably contribute to the overall heat transfer process. Thus, the proposed models have exhibited a high accuracy as compared to solvers that utilize finite element methods, while the models can be employed for fast and reasonably thermal analysis of 3-D systems. Beyond these models, thermal management policies for MPSoCs either planar or vertically integrated have been implemented. The development of these policies was performed in collaboration with ETHZ. The outcome is a technique that properly adjusts the operating frequency of the cores such that specific thermal constraints are obeyed. In addition, another method for the placement of therma sensors that provide the temperature information for each of the cores has been realized.

-- The above is new material, not present in the Y2 deliverable --

2.2 Individual Publications Resulting from these Achievements

ETHZ

1. [SCT10] Nikolay Stoimenov, Samarjit Chakraborty, Lothar Thiele: An Interface Algebra for Estimating Worst-Case Traversal Times in Component Networks. Leveraging Applications of Formal Methods, Verification, and Validation, 4th International Symposium on Leveraging Applications, ISoLA 2010, Lecture Notes in Computer Science 6415, Springer, Heraklion, Crete, Greece, pages 198-213, October, 2010.

2. [SPCTC10] Andreas Schranzhofer, Rodolfo Pellizzoni, Jian-Jia Chen, Lothar Thiele, Marco Caccamo: Worst-Case Response Time Analysis of Resource Access Models in Multi-Core Systems. Proceedings of the 47th Design Automation Conference (DAC), ACM, Anaheim, California, USA, pages 332--337, June, 2010.

3. [LDT10] Kai Lampka, Simon Perathoner, Lothar Thiele: Analytic Real-Time Analysis and Timed Automata: A Hybrid Methodology for the Performance Analysis of Embedded Real-Time Systems. Design Automation for Embedded Systems, Springer Science+Business Media, LLC, Vol. 14, No. 3, pages 193-227, 2010.

TU Braunschweig

4. [SNE10] Simon Schliecker, Mircea Negrean and Rolf Ernst, **"Bounding the Shared Resource Load for the Performance Analysis of Multiprocessor Systems,"** in *Proc. of Design, Automation, and Test in Europe (DATE)*, (Dresden, Germany), March 2010

5. [RE10] Jonas Rox and Rolf Ernst, **"Formal Timing Analysis of Full Duplex Switched Based Ethernet Network Architectures,"** in *SAE World Congress*, vol. System Level

Architecture Design Tools and Methods (AE318), (Detroit, MI, USA), SAE International, April 2010.

6. [NSE10] Mircea Negrean, Simon Schliecker and Rolf Ernst, "**Timing Implications of Sharing Resources in Multicore Real-Time Automotive Systems**," in *SAE World Congress*, vol. System Level Architecture Design Tools and Methods (AE318), (Detroit, MI, USA), April 2010 *(this paper was selected for SAE Journals – see next publication)*
7. [NSE+10] Mircea Negrean, Simon Schliecker, and Rolf Ernst, **"Timing Implications of Sharing Resources in Multicore Real-Time Automotive Systems,"** *SAE International Journal of Passenger Cars - Electronic and Electrical Systems*, vol. 3, No. 1, pp. 27-40, August 2010 *(previous publication was selected for SAE Journals)*

Linköping

8. [BAE11] Min Bao, Alexandru Andrei, Petru Eles, Zebo Peng, On-Line Temperature-Aware Idle Time Distribution for Leakage Energy Optimization, Proceedings of International Symposium on Electronic Design, Test and Applications (DELTA 2011), New Zealand, January 2011.
9. [BAE10] Min Bao, Alexandru Andrei, Petru Eles, Zebo Peng, Temperature-Aware Idle Time Distribution for Energy Optimization with Dynamic Voltage Scaling, Proceedings of Design Automation and Test in Europe (DATE'2010), Dresden, Germany, March 8-12, 2010.

DTU

10. Mikkel Koefoed Jakobsen, Jan Madsen, and Michael R. Hansen: DEHAR: a Distributed Energy Harvesting Aware Routing Algorithm for Ad-hoc Multi-hop Wireless Sensor Networks. In proc. of IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM 2010)
11. Tolga Ovatman, Aske W. Brekling, Michael R. Hansen, Cost Analysis for Embedded Systems: Experiments with Priced Timed Automata. Electronic Notes in Theoretical Computer Science 238 (2010) 81ñ95
12. Jan Madsen, Michael R. Hansen, Aske Brekling: Modelling and Analysis for Embedded Systems. Chapter in Model-Based Design for Embedded Systems, CRC Press, Pages 121-145, 2010
13. Michael R. Hansen. Efficient Model Checking for a Hybrid Duration Calculus. In proc. of International Workshop on Hybrid Logic and Applications (HyLo 2010).
14. Aske Brekling. Modelling and Analyses of Embedded Systems Design. PhD Thesis. Informatics and Mathematical Modelling, Technical University of Denmark, Report no. IMM-PHD-2010-236, 2010

UNIBO

15. [LMRB10] Michele Lombardi, Michela Milano, Martino Ruggiero, Luca Benini: Stochastic allocation and scheduling for conditional task graphs in multi-processor systems-on-chip. J. Scheduling (SCHEDULING) 13(4):315-345 (2010)

KTH

16. Fahimeh Jafari, Zhonghai Lu, Axel Jantsch, and Mohammad Hossein Yaghmaee. Buffer optimization in network-on-chip through flow regulation. IEEE Transactions on Computer Aided Design (TCAD), 29(12):1973-1986, December 2010

17. Xiaowen Chen, Zhonghai Lu, Shuming Chen, and Axel Jantsch. Run-time partitioning of hybrid distributed shared memory on multi-core network-on-chips. In The 3rd IEEE International Symposium on Parallel Architectures, Algorithms and Programming (PAAP 2010), Dalian, China, December 2010.

18. Xiaowen Chen, Shuming Chen, Zhonghai Lu, and Axel Jantsch. Area and performance optimization of barrier synchronization on multi-core network-on-chips. In 3rd IEEE International Conference on Computer and Electrical Engineering (ICCEE), Chengdu, China, November 2010.

19. Xiaowen Chen, Shuming Chen, Zhonghai Lu, and Axel Jantsch. Multi-FPGA implementation of a network-on-chip based many-core architecture with fast barrier synchronization mechanism. In Proceedings of the IEEE Norchip Conference, Tampere, Finland, November 2010.

20. Chaochao Feng, Zhonghai Lu, Axel Jantsch, Jinwen Li, and Minxuan Zhang. A reconfigurable fault-tolerant deflection routing algorithm based on reinforcement learning for networks-on-chip. In Proceedings of the International Workshop on Network on Chip Architectures (NoCArc), Atlanta, Gorgia, November 2010.

21. Abbas Eslami Kiasari, Axel Jantsch, and Zhonghai Lu. A framework for designing congestion-aware deterministic routing. In Proceedings of the International Workshop on Network on Chip Architectures (NoCArc), Atlanta, Gorgia, November 2010.

22. Xiaowen Chen, Zhonghai Lu, Axel Jantsch, and Shuming Chen. Handling shared variable synchronization in multi-core network-on-chip with distributed memory. In International SOC Conference, Las Vegas, Nevada, September 2010.

23. Chaochao Feng, Zhonghai Lu, Axel Jantsch, Jinwen Li, and Minxuan Zhang. FoN: Fault-on-neighbor aware routing algorithm for networks-on-chip. In International SOC Conference, Las Vegas, Nevada, September 2010.

24. Zhipeng Chen and Axel Jantsch. A worst case performance model for TDM virtual circuit in nocs. In Proceedings of the International Workshop on Network on Chip, Zheng Zhou, China, September 2010.

25. Jun Zhu, Ingo Sander, and Axel Jantsch. HetMoC: Heterogeneous modeling in systemc. In Proceedings of the Forum on Design Langauges (FDL), Southhampton, UK, September 2010.

26. Xiaowen Chen, Zhonghai Lu, Axel Jantsch, and Shuming Chen. Supporting efficient synchronization in multi-core NoCs using dynamic buffer allocation technique. In Proceedings of the IEEE Annual Symposium on VLSI, Kefalonia, Greece, July 2010.

27. Ming Liu, Zhonghai Lu, Wolfgang Kuehn, and Axel Jantsch. Inter-process communication using pipes in FPGA-based adaptive computing. In Proceedings of the IEEE Annual Symposium on VLSI, Kefalonia, Greece, July 2010.

28. Abdul Naeem, Xiaowen Chen, Zhonghai Lu, and Axel Jantsch. Scalability of weak consistency in NoC based multicore architectures. In Proceedings of the IEEE International Symposium on Circuits and Systems (ISCAS), Paris, France, June 2010.

29. Ming Liu, Zhonghai Lu, Wolfgang Kuehn, and Axel Jantsch. Reducing FPGA reconfiguration time overhead using virtual configurations. In Proceedings of the 5th International Workshop on Reconfigurable Communication Centric Systems-on-Chip, Karlsruhe, Germany, May 2010.

30. Amr Helmy, Laurence Pierre, and Axel Jantsch. Theorem proving techniques for formal verification of noc communications with non-minimal adaptive routing. In Proceedings of

the 13th IEEE International Symposium on Design & Diagnostics of Electronic Circuits & Systems, Vienna, Austria, April 2010.

31. Xiaowen Chen, Zhonghai Lu, Axel Jantsch, and Shuming Chen. Supporting distributed shared memory on multi-core network-on-chips using a dual microcoded controller. In Proceedings of the confernece for Design Automation and Test in Europe, Dresden, Germany, March 2010.
32. Jun Zhu, Ingo Sander, and Axel Jantsch. Pareto efficient design for reconfigurable streaming applications on CPU/FPGAs. In Proceedings of Design Automation and Test in Europe (DATE '10), Dresden, Germany, March 2010.
33. Fahimeh Jafari, Zhonghai Lu, Axel Jantsch, and Mohammad H. Yaghmaee. Optimal regulation of traffic flows in networks-on-chip. In Proceedings of the Design Automation and Test Europe Conference (DATE), Dresden, March 2010.
34. Ming Liu, Zhonghai Lu, Wolfgang Kuehn, and Axel Jantsch. FPGA-based adaptive computing for correlated multi-stream processing. In Proceedings of the Conference Design, Automation and Test Europe, Dresden, Germany, March 2010.
35. Jun Zhu, Ingo Sander, and Axel Jantsch. Performance analysis of reconfigurations in adaptive real-time streaming applications. ACM Transactions in Embedded Computing Systems - Special issue on Embedded Systems for Real-time Multimedia, 2010.
36. Jun Zhu, Ingo Sander, and Axel Jantsch. Constrained global scheduling of streaming applications on MPSoCs. In Proceedings of the conference on Asia South Pacific Design Automation (ASP-DAC '10), Taipei, Republic of China, January 2010.

CEA LIST

37. T. Sassolas, N. Ventroux,G. Blanc. A Power-Aware Online Scheduling Algorithm for Streaming Applications in Embedded MPSoC. 20th International Workshop on Power and Timing Modeling, Optimization and Simulation (PATMOS), 2010 , Grenoble , France

EPFL

38. [SMBM10] C. Seiculescu, S. Murali, L. Benini, and G. De Micheli, "SunFloor 3D: A Tool for Networks on Chip Topology Synthesis for 3D Systems on Chips," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, Vol. 29, No. 12, pp. 1987-2000, December 2010.
39. [SMBM10b] C. Seiculescu, S. Murali, L. Benini, and G. De Micheli, "Comparative Analysis of NoCs for Two-Dimensional Versus Three-Dimensional SoCs Supporting Multiple Voltage and Frequency Islands," *IEEE Transactions on Circuits and Systems II: Express Briefs*, Vol. 57, No. 5, pp. 364 - 368, May 2010.
40. [TPM10] I. Tsioutsios, V. F. Pavlidis, and G. De Micheli, "Physical Design Tradeoffs in Power Distribution for 3-D ICs," *Proceedings of the IEEE International Conference on Electronic Circuits and Systems*, pp. 435-438, December 2010.
41. [ZAMB] F. Zanini, D. Atienza, G. De Micheli, and S. P. Boyd, "Online Convex Optimization-Based Algorithm for Thermal Management of MPSoCs," in *Proceedgins of the ACM Great Lakes Symposium*, pp. 203-208, May 2010.
42. [PXTM10] V. F. Pavlidis, H. Xu, I. Tsioutsios, and G. De Micheli, "Synchronization and Power Integrity Issues in 3-D ICs," *Proceedings of the IEEE Asia Pacific Conference on Circuits and Systems,* pp. 536-539, December 2010.
43. [BCTBPM10] S. Bobba, A. Chakraborty, O. Thomas, P. Batude, V. F. Pavlidis, and G. De Micheli, "Performance Analysis of 3-D Monolithic Integrated Circuits," *Proceedings of the IEEE 3-D System Integration Conference,* November 2010.
44. [XPM10] H. Xu, V. F. Pavlidis, and G. De Micheli, "Process-Induced Skew Variations for Scaled 2-D and 3-D ICs," *Proceedings of the ACM/IEEE International Workshop on System Level Interconnect Prediction,* pp. 17-24, June 2010.

-- The above are new references, not present in the Y2 deliverable --

2.3 Interaction and Building Excellence between Partners

Strategies to Control Environmentally Powered Systems (ETHZ, UNIBO)
The interaction in this activity has been between ETHZ and UNIBO.

- Joint work in developing new algorithms to perform application control inenvironmentally powered systems.
- As a result, a joint journal paper has been published [MTBB10].

Unifying approaches for compelx stream descriptions (ETHZ, TUBS)
The interaction in this activity has been between ETHZ and TUBS.

- Joint work in developing new schemes to analyse hierarchical event streams.
- The MPA and Symta/S tool coupling has been improved in order to include the new class of hierarchical event streams.
- A joint publication has been written [PRTLR10].

Modelling of Shared Resources in Multiprocessor Systems (TUBS, ETHZ)

The interaction in this activity has been between TUBS and ETHZ.

- Face-to-face meeting in Brussel: Andreas Schranzhofer (ETHZ), and Simon Schliecker and Mircea Negrean (TUBS) discussed open issues and identified possible improvements on the modelling and analysis approaches of multi-core systems with shared resources. Results were presented at the ArtistDesign cluster meeting in Leuven, Belgium.

Runtime layer design for many-cores architectures (CEA LIST, UNIBO)

Interaction between CEA LIST and University of Bologna in this activity is based on face-to-face meetings in Saclay and Genoble and numerous phone meetings to clarify technical aspects regarding Power Management and fine grain parallelism management in MPSoC.

Modeling and analysis for QoS optimisation of real-time control applications (Linkoping, Lund)

The interaction in this activity has been between Linköping and Lund.

- Joint work to develop control performance models, scheduling, and optimisation algorithms.
- Soheil Samii and Anton Cervin have done several visits at Lund and Linköping, respectively.
- Two joint papers have been published.

Modeling and analysis of heterogeneous systems (KTH, DTU)

The interaction in this activity has been between KTH and DTU, and has been supported by the Artemis project SYSMODEL.

- Joint work to develop the System Functionality Framework (SFF) based on the ForSyDe model.
- Joint development of SystemC templates to ease the creation of SFF models for SMEs.
- Mikkel Koefoed Jakobsen and Seyed Hosein Attarzadeh Niaki have done several visits at KTH and DTU respectively.
- Two joint 1-day workshops have been held.
- A first joint paper is in the writing.

Modeling and Verification of Embedded Systems (DTU, AAU)

The interaction in this activity has been between DTU and AAU.

- The timed automaton formulation of the ARTS model using UPPAAL has been done by DTU.
- Joint discussions and meetings have been done and research version of UPPAAL supporting features needed for efficient modelling of the ARTS model has been provided by AAU.
- DTU and AAU has each provided a chapter (with cross references) on analysis and scheduling in the book, “Model-Based Design of Heterogeneous Embedded Systems” by Gabriela Nicolescu, Pieter J. Mosterman.

Formal verification of design properties of hardware architectures (DTU, Virginia Tech)

The interaction in this activity has been between DTU and Virginia Tech.

- Skype discussions on edicational issues of using Gezel for teaching embedded systems design.

A framework for automatic parallelization, static and dynamic memory optimization in MPSoC platforms

- NTUA and Imec closely collaborated on this topic. This is reflected by a joint publication on the subject during the intership of Yiannis Iosifidis (NTUA) at Imec working in the context of the NMEMEE and 2PARMA European projects.

Synthesis Tool for 3-D NoC and Analysis of 2- *versus* 3-D NoC topologies with power management schemes (UNIBO, EPFL)

The interaction in this activity has been between EPFL and UNIBO

- There has been active collaboration during the development of the tool
- Meetings and visits between UNIBO and EPFL have been made to determine the specifics of the tool as well as to elaborate on the DVFS schemes that were assumed in the comparison of the different NoC topologies

-- Changes wrt Y2 deliverable --

This is new text reflecting the interactions of Y3.

2.4 Joint Publications Resulting from these Achievements

1. [BCTBG10] Andrea Bartolini, Matteo Cacciari, Andrea Tilli, Luca Benini, and Matthias Gries. “A virtual platform environment for exploring power, thermal and reliability management control strategies in high-performance multicores.” (GLSVLSI '10)
2. [PRTLR10] Simon Perathoner, Tobias Rein, Lothar Thiele, Kai Lampka, Jonas Rox: Modeling Structured Event Streams in System Level Performance Analysis. ACM SIGPLAN/SIGBED Conference on Languages, Compilers and Tools for Embedded Systems (LCTES), ACM, Stockholm, Sweden, pages 37-46, April, 2010.
3. [MTBB10] Moser, C.; Thiele, L.; Brunelli, D.; Benini, L.: Adaptive Power Management for Environmentally Powered Systems. IEEE Transactions on Computers (59:4), 2011.
4. [CASA10] Bernard Candaele, Sylvain Aguirre, Michel Sarlotte, Iraklis Anagnostopoulos, Sotirios Xydis, Alexandros Bartzas, Dimitris Bekiaris, Dimitrios Soudris, Zhonghai Lu, Xiaowen Chen, Jean-Michel Chabloz, Ahmed Hemani, Axel Jantsch, Geert Vanmeerbeeck, Jari Kreku, Kari Tiensyrja, Fragkiskos Ieromnimon, Dimitrios Kritharidis, Andreas Wiefrink, Bart Vanthournout, and Philippe Martin. Mapping optimisation for scalable multi-core architecture: The MOSART approach. In Proceedings of the IEEE Annual Symposium on VLSI, Kefalonia, Greece, July 2010.
5. [OCLDB10] M. Ojail, K. Ben Chehida, Y. Lhuillier, R. David, L. Benini. Synchronous Reactive Task Management for Fine Grain Parallelism in Manycores Architectures. Submitted to Workshop on Parallel Programming and Run-time Management Techniques for Many-core Architectures 2010, Hannover. *This paper is also quoted in MPSOC architecture report since it involves HW ans SW activities of both clusters.*
6. [SEP11] Soheil Samii, Petru Eles, Zebo Peng, Anton Cervin, Design Optimization and Synthesis of FlexRay Parameters for Embedded Control Applications, Proceedings of International Symposium on Electronic Design, Test and Applications (DELTA 2011), New Zealand, January 2011.
7. [SEP10] Soheil Samii, Petru Eles, Zebo Peng, Paulo Tabuada, Anton Cervin, Dynamic Scheduling and Control-Quality Optimization of Self-Triggered Control Applications, Proceedings of 31st IEEE Real-Time Systems Symposium (RTSS10), San Diego, CA, USA, November 30-December 3, 2010.
8. [SMBM10] C. Seiculescu, S. Murali, L. Benini, and G. De Micheli, “SunFloor 3D: A Tool for Networks on Chip Topology Synthesis for 3D Systems on Chips,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, Vol. 29, No. 12, pp. 1987-2000, December 2010.
9. [SMBM10b] C. Seiculescu, S. Murali, L. Benini, and G. De Micheli, “Comparative Analysis of NoCs for Two-Dimensional Versus Three-Dimensional SoCs Supporting Multiple Voltage and Frequency Islands,” *IEEE Transactions on Circuits and Systems II: Express Briefs*, Vol. 57, No. 5, pp. 364 - 368, May 2010.
10. [SMBM10c] C. Seiculescu, S. Murali, L. Benini, and G. De Micheli, “A Method to Remove Deadlocks in Networks-on-Chips with Wormhole Flow Control,” *Proceedings of the Design Automation and Testing Conference*, pp. 1625-1628, March 2010.
11. [ZAJM] F. Zanini, D. Atienza, C. N. Jones, and G. De Micheli, “Temperature Sensor Placement in Thermal Management Systems for MPSoCs,” *Proceedings of the IEEE International Conference on Circuits and Systems*, pp. 1065-1068, May 2010.
12. 8. [IMM10] Y. Isofidis, A. Mallik, S. Mamagkakis, E. De Greef, A. Bartzas, D. Soudris and F. Catthoor, “ A framework for automatic parallelization, static and dynamic memory optimization in MPSoC platforms”, in Proceeding of the 47th Design Automation Conference (DAC), 2010.

-- The above are new references, not present in the Y2 deliverable --

2.5 Keynotes, Workshops, Tutorials

Keynote : P2012: A many-core platform for 10Gops/mm2 multimedia computing (Luca Benini, UNIBO)
21st IEEE International Symposium on Rapid System Prototyping (RSP)
Fairfax, Virginia, USA - June 8-11, 2010

The IEEE International Symposium on Rapid System Prototyping (RSP) explores trends in Rapid Product Development of Computer Based Systems. Its scope ranges from formal methods for the verification of software and hardware systems to case studies of actual software and hardware systems. It aims to bring together researchers from the hardware and software communities to share their experiences and to foster collaboration of new and innovative Science and Technology.

http://www.rsp-symposium.org/rsp2010/index.html

Keynote : Many-core platforms for embedded computing: 2012 and beyond (Luca Benini, UNIBO)
S02-XXI Jornadas de Paralelismo, JP2010 (SARTECO) - CEDI Congeso Espanol di Informatica
Valencia, SPAIN– 7-10 September 2010

Abstract :*Programmability is a key requirement for fast time-to-market and agile adaptation to rapidly evolving multimedia standards and customer expectations. Unfortunately, programmable architectures come with order-of-magnitude computational density and energy efficiency gaps with respect to custom-fit hardware. Is there a way to escape the flexibility vs. efficiency dualism? Is nano-scale silicon technology adding new facets to this "no free lunch" view? In this talk I will describe the architectural foundations of STMicroelectronics Platform 2012 project and provide some insight on how we hope to give positive answers to these fundamental questions*

http://cedi2005.ugr.es/2010/contenido.php?apartado=organizacion&sub=presentacion/

Tutorial: Modeling and Analyzing Real-Time Mutlriprocessor Systems
ESWEEK 2010
Scottsdale, USA – 24.10.2010
The presentation introduced the component-based approach that has been developed in ArtistDesign as a joint activity of UBS and ETHZ. The presentation hase been given by Lothar Thiele (ETHZ).

Workshop: Temperature Aware 3D MPSoC
Lausanne, Switzerland – 26.11.2010
The workshop has been organized by David Atienza (EPFL). The porpuse was to discuss various approaches to analyze the temperature of complex 3D MPSoC systems. The focus of this event was on the comparison of various levels of abstraction and the corresponding analysis models and methods.

Tutorial: Performance Analysis of Distributed Embedded Systems
*Rabat, Morocco – 11.7. – 16.7.*2010
The tutorial was part of the ArtistDesign School in Morocco. The presentation introduced the component-based approach that has been developed in ArtistDesign as a joint activity of UBS and ETHZ. The presentation hase been given by Lothar Thiele (ETHZ).

Invited Lecture: Formal Performance Analysis and Optimization of Safety-related Embedded Systems
(Rolf Ernst, TU Braunschweig)
Artist Summer School Europe 2010
Autrans, France, September 6, 2010

The lecture was given at the Artist Summer School Europe organized by the Artist Design European Network of Excellence on Embedded Systems Design. An important part of the lecture was dedicated to formal performance analysis in general, and with a focus on required analysis extensions that allow the application of performance analysis methods to safety related system design. *(Further topics addressed in this lecture are related to the activity reports on Platform and MPSoC Design and on the Design for Adaptivity Activity)*

http://www.artist-embedded.org/artist/Invited-Speakers,2065

Invited Talk: Providing Real-Time Guarantees on Multi-Core Processors
(Simon Schliecker/ Mircea Negrean, TU Braunschweig)
Symtavision NewsConference 2010
Braunschweig, Germany – September 29 -39, 2010

The SymTA/S NewsConference is an annual event organized by the Symtavision GmbH that brings together engineers, managers, technology experts and researchers in the field of embedded real-time systems. Compared to previous years, the 4th edition of the Symtavision NewsConference has been expanded by a technical day, with parallel practice and research tracks. TU Braunschweig was invited to present current research results on real-time analysis methods for multi-core systems.

http://www.symtavision.com/newsconference2010.html

Tutorial: Modeling and Analyzing Real-Time Multiprocessor Systems
(Simon Schliecker, TU Braunschweig, Lothar Thiele, ETHZ and other participants)
Embedded Systems Week (ESWeek) 2010
Scottsdale AZ, USA, October 24 -29, 2010

Embedded Systems Week is an exciting event which brings together conferences, tutorials, and workshops centred on various aspects of embedded systems research and development. The tutorial at ESWeek 2010 was given by Lothar Thiele (ETHZ), Simon Schliecker (TU Braunschweig), Marco Bekooij (NXP semiconductors), Maarten Wiggers (University of Twente) and Edward A. Lee (University of California Berkeley). This tutorial presented an overview and positioning of four recently proposed approaches for timing constraints verification in real-time multiprocessor systems.

http://www.esweek.org/

Invited seminar: Promises and limitations of 3-D integration (Axel Jantsch, Matthew Grange, and Dinesh Pamunuwa)

Natinal University of, Defense Technology
Changsha, China, December 2010

Invited seminar: Memory architecture and management in a NoC platform (Axel Jantsch, Xiaowen Chen, Abdul Naeem, Yuang Zhang, Sandro Penolazzi, and Zhonghai Lu)
Fudan University, Shanghai, China*, December 2010*
Nanjing University, China, *December 2010*

Invited talk**: Network on Chip Technology for Telecom Applications** (Axel Jantsch)
Fudan-Huawei Workshop, Shanghai, China*, December 2010*

Invited lecture**: Predictable communication performance in on-chip networks** (Axel Jantsch)
University of Turku, Finland, December 2010

Keynote : R. DAVID, CEA LIST, Low Power management in embedded multi-core architectures
MPSOC
Gifu, Japan, july 2010

Abstract :Due to the complexity increase of embedded applications, multi-core systems on chip (MPSOC) are becoming the mainstream for architecture design. Indeed, according to ITRS, the number of cores in high-end systems will exceed 100 cores in 2012. With an expected 32% a year increase in the number of cores per die, the concept of multi-cores will even evolve to many-cores in the coming years. Managing complexity in such system is a challenge that cannot be talked only by off-line application analysis and compilation techniques. In fact to deal with variability coming from technology as well as advance embedded application that are highly data-dependant, runtime resources management strategies have to be put in place. This talk deals with this new challenge of dynamically managing tens to hundreds of core. In particular this talk will focus on the power management of such devices.

Seminar: Energy Harvesting (Jan Madsen and Michael R. Hansen, DTU)
Special Interest Group on Green-IT, Infinit innovation network on ICT
*Lyngby, Denmark – 11.6.*2010
DTU organized the SIG seminar on energy harvesting. The event had 55 participants, where half were from industry.
http://www.infinit.dk/dk/nyheder_og_arrangementer/arrangementer/afsluttede_aktiviteter/20100531_01.htm

Seminar: Safety-Critical Systems (Paul Pop, DTU)
Special Interest Group on Safety-Critical Systems, Infinit innovation network on ICT
*Lyngby, Denmark – 12.2.*2010
Embedded systems are increasingly used in safety-critcal application areas, such as medical, railway, military, aerospace and factory systems, where a failure can endanger human life or the environment. Safety-critical embedded systems are becoming more complex, and use software and hardware to implement part of their functions. This is the reason why the national innovation network InfinIT is starting an interest group on safety-critical systems. The objective of the seminar is to present the current challenges in developing safety-critical systems and to identify the focus of the interest group. The event had 40 participants, where more than half were from industry.
http://www.infinit.dk/dk/nyheder_og_arrangementer/arrangementer/afsluttede_aktiviteter/20100125_01.htm

Course: Automated Formal Methods for Embedded Systems (Jan Madsen and Michael R. Hansen, DTU)
ARTIST Graduate Course
Lyngby, Denmark – June 14-28 2010
DTU organized the ARTIST graduate course on automated formal methods for embedded systems. The event had 12 participants. The topics covered were, model-based development and validation of multi-robot cooperative systems, Simulation of Networked Embedded Control Systems Using TrueTime, logical approach to modelling and analysis of resource constraints, and specification and verification with proof scores in CafeOBJ.

Seminar: Model-Based Design (Jan Madsen and Michael R. Hansen, DTU)
Company workshop at KK-Electronics
*Ikast, Denmark – 8.12.*2010
DTU organized a company seminar for 22 software engineers at KK-Electronics. The aim was to give an overview of model-based design of embedded systems. The company develops windturbine controllers for Siemens Windpower, Vestas and other manufactures of windmills.

Seminar: Embedded/Safety Critical Systems (Jan Madsen, DTU)
IEEE Denmark
Copenhagen, Denmark – 1.11.2010
For some years embedded systems for safety critical applications have increasingly taken over important functions in our working and everyday life. One reason for this to happen is obviously the ability of such systems to relieve humans from the burden of constant monitoring processes in detail and taking corrective actions if needed. Moreover, the potential of such systems to eliminate human errors is getting increasingly important as keeping track and overview of the many variables encountered in the evermore complex environment of the modern society is difficult for a human mind, if not impossible. At the same time, real life examples show system design failures that have caused safety critical systems to respond with unintended actions, in case a system state is allowed appear in unforeseen situations. Introducing the seminar with a lecture on Embedded Systems, Professor Jan Madsen, IMM/DTU, is also addressing this dilemma along with the research on Embedded Systems carried out at his Institute. To illustrate the applications of embedded systems in an engineering context, IDA and IEEE Denmark has succeeded in getting acceptance of speakers from industry within the fields of Space, Aviation, Railways and Medicine.
http://meetings.vtools.ieee.org/meeting_view/list_meeting/2963

Workshop: SYSMODEL Modeling Methodologies (Paul Pop, DTU, Ingo Sander, KTH)
SYSMODEL
Kista, Stockholm – May 27-28.2010
Workshop on modeling methodologies for the ForSyDe modeling framework. Identification and planning of industrial use cases, to evaluate the capabilities of the modeling framework.

Workshop: ForSyDe Modeling Workshop
SYSMODEL
Kista, Stockholm – Feb 1-2.2010
Joint workshop between KTH and DTU on the development of the ForSyDe modelling framework.

Workshop: ASAM Design Flow (Jan Madsen, DTU)
ASAM
Cagliari, Sardinia – Sept. 23-24.2010
DTU and TUBS participated in the 2-day ASAM workshop aimed at defining the design flow and tool integration for the ASAM project.

Keynote : De Micheli, Giovanni: *Nanosystems: devices, circuits, architectures and applications*.
Conference name International Symposium on Circuits and Systems (ISCAS)

May 30 - June 2nd, 2010.

The IEEE International Symposium on Circuits and Systems (ISCAS) is the world's premier networking forum of leading researchers in the highly active fields of theory, design and implementation of circuits and systems.

The Symposium will focus on circuits and systems employing nanodevices (both extremely scaled CMOS and non-CMOS devices) and circuit fabrics (mixture of standard CMOS and evolving nano-structure elements) and their implementation cost, switching speed, energy efficiency, and reliability.

http://www.iscas2010.org/index.php?option=com_content&view=section&id=31&Itemid=131#keynote1

Keynote : De Micheli, Giovanni: Nanosystems for a healthier and safer tomorrow
Conference name Symposium on Transformational Information Engineering and Science

Singapore, Nanyang Technological University, January 28-29, 2010

This symposium brings together diverse and notable experts to discuss the technical and societal implications of the ubiquitous systems people now rely upon -- billions of times each day -- to transform information from one domain to another. These "transformations" occur every time a cell phone converts electronic signals into sound, a shopper encrypts a credit card number online or a moviegoer's eye converts moving pictures into information the brain can process.

http://www.ntu.edu.sg/ISNE/event/TIES/Pages/ProgrammeDetails.aspx

-- The above is new material, not present in the Y2 deliverable --

3. Milestones, and Future Evolution

3.1 Problems to be Tackled in Year 4 *(Jan 2011 – Dec 2011)*

UNIBO will continue on the research activity about scheduling problem on multicore systems. We expected futrther improvements on refining Scheduling algorithms. The goal is to ensure the highest utilization of the multicore system and to develop efficient scheduling of multi-task applications with inter-task dependencies. Moreover UNIBO will continue the research activity on power, thermal and reliability management control strategies in collaborations with Intel Labs.

In year 4, **ETHZ** will focus its joint activities related to the analysis of multicore systems. In particular, we will concentrate on the interference of task executions on joint resources such as memory and buses. In addition, we will investigate together with partners at **EPFL** the analysis of temperatures in complex 3D MPSoC systems.

During the fourth year, **TU Braunschweig** will continue to work on providing real-time analysis methods for multiprocessor systems with shared resources. The goal is to extend existing solutions and to develop new analyses that suit new and more complex multiprocessor applications. Similarities in the approaches provided by **TU Braunschweig** and **ETHZ** have been identified in the last year. The joint work will continue in order to identify solutions on open issues on the analysis approaches for multiprocessor systems.

Multi-core systems with shared resources considered in the activity of previous years require new analysis approaches in case they can switch between different operational modes. **TU Braunschweig** will work on analysis approaches for multi-mode systems.

In year 4, **CEA LIST** and **UNIBO** will tackle the challenge of managing resources of manycores architecture by exending actual framework. Main challenge is to keep the ability to completely hide the hardware complexity to the firmware developers despite a clusterized architecture with a non uniform memory space.

During the fourth year, the Linköping group will work on the extension of the elaborated temperature models to multicore systems.

During the fourth year the groups at **Linköping** and **DTU** will continue their work on fault tolerant distributed systems. The focus will be on fault modelling of the communication infrastructure aiming at the synthesis of fault tolerant communication over FlexRay buses.

DTU will continue its efforts on analytical performance models for both system-level, such as ARTS and Duration Calculus, and low-level, such as a formalized version of the Gezel hardware description language. This will be done in cooperation with Oldenburg, AAU and Virginia Tech.

KTH and **DTU** will continue their work on a SystemC based modeling framework based on the ForSyDe formal model. The framework will be developed to support untimed, synchronous time, discrete time and continuous time models of computation. The framework will be evaluated on a number of use cases defined by the SMEs in the SYSMODEL project.

KTH will continue its work on performance analysis and dimensioning of MPSoC communication infrastructure. So far the focus has been on *static, design-time* solutions. In year 4 KTH will focus on deploying and extending the developed techniques to *dynamic, run-time* problems and their solutions. The underlying techniques are based on worst case analysis with Network Calculu, performance contracting, and traffic shaping. Run-time traffic shaping has to potential to adaptively re-allocate communication resources in order to provide worst case performance guarantees with minimal resource usage.

As a new activity, KTH is initiating work on fault-tolerant and reliable MPSoC communication architectures. Further technology scaling will increase the need for an ability to tolerate faults and errors. KTH will focus on techniques to identify, correct, and tolerate faults on the link layer, the network layer and the transport layer of the MPSoC communication fabric.

EPFL will focus their work on 3-D integrated systems on be the global infrastructure of these systems. The choice of these research problems for Year 4 is rationalized, since clock and power distribution are expected to be predominant problems in the 3-D design process that need to be addressed. One of the main objectives is to provide accurate models that enhance the timing analysis of these complex systems. These models will consider the manufacturing variability, which poses different constraints due to the vertical integration in the timing analysis procedure.

The knowledge acquired by **EPFL** on their concurrent activities related to the global infrastructure of 3-D systems will be adapted to analyse a microprocessor-DRAM system in terms of power integrity. This tool will be enhanced with thermal models in order to consider the effects of temperature on the behaviour of the current loads (i.e., the current requirements of the circuits as a function of the temperature) and the increased interconnect resistance due to the elevated temperature.

-- Changes wrt Y2 deliverable --

Most of this is new text. It is an update based on the descripton from year 2.

3.2 Current and Future Milestones

Integration of Symta/S - MPA and unifying approaches for hierarchical scheduling

More recent versions of SymTA/S employ more general event models (denoted here as delta-functions) which can be considered as the pseudo-inverses of arrival curves. The conversion of arrival curves to delta-functions has been completed. The conversion of delta-functions to arrival curves is still open and will be tackled in the next year.

This milestone has been fulfilled. When converting event models from the SymTA/S to the MPA representation an issue with the interface SymTA/S → MPA is that it is not always tight, i.e., it can introduce pessimism for the analysis results. ETHZ and TUBS invested efforts to find a solution for reducing the pessimism. The conclusion is that a lossless interface for converting event models from SymTA/S to MPA seems not realizable. It is, however, possible to control the precision of the conversion.

In Year 4 the tool coupling with the ability of adjusting the conversion precision will be used to study the trade-off between accuracy and runtime.

Performance analysis of inter-task synchronization in multiprocessor systems (TU Braunschweig)

In the 3rd year, the work on unifying the methodologies to capture shared resource synchronization and shared memory accesses shall be continued. The framework shall be applied to new applications.

This milestone has been fulfilled. A method that captures more accurately the interference between different cores that share common resources has been published in [SNE10]. The framework has been applied to investigate the impact of the shared resource usage on the systems' timing for different multiprocessor design options

[NSE+10]. Further work was performed to investigate the applicability of the analysis methods to particular problems in the industry, e.g. in the automotive domain.

In the 4th year, the work on the methodology to capture the shared resource synchronization in multiprocessor systems will continue. The framework developed in the previous years will be extended to consider new system setups. Further on, effort will be devoted to provide analysis approaches for multi-mode systems.

Hybrid approach combining Real-Time Analysis and Timed Automata

Concerning the hybrid approach combining Real-Time Analysis (RTC) and Timed Automata, CEA will try to extend the schedulability analysis of RTC to state-based schedulers by applying the event generator approach. For example, we are working on the feasibility analysis for adaptive DVS scheduling, to optimistically minimize the energy when the system is lightly loaded by executing at low speeds, and to pessimistically meet the timing constraints when the system is heavily loaded by executing at the maximum speed of the system. ETHZ will extend its framework for coupling timed automata with MPA and attempt to improve its scalability to large distributed embedded systems.

In the third year, we have been able to couple the RTC Analysis framework of ETHZ with state-based performance analysis methods. Therefore, this milestone has been reached, see [LDT10].

The focus in the fourth year will be the extension of this hybrid approach to the analysis of bus-based multicore systems, i.e. systems with several interacting resources. New methods of interference analysis of task executions on joint resources developed in cooperation with other partners in this activity.

Contract based architecture dimensioning

Based on the definition of performance contracts between IPs and the SoC infrastructure (done in year 1), in year 2 KTH will work on formulating the problem of dimensioning the infrastructure, given a set of contracted flows and proposing methods for solving it.
The above milestone has been fulfilled partially. The regulation spectrum has been defined, which defines the space of possible contract parameters and the the optimization opportunities for traffic shaping. An optimization algorithm has been developed and implemented and will hopefully be published in year 3.

During the third year we will further study different aspects of infrastructure dimensioning based on conracted flows, and develop optimization methods and algorithms**.**

The above milestone has been fulfilled. The basic concepts have been built and a static resource dimensioning method has been developed.

The focus in year 4 will be to to formulate the problem for dynamic renegotiation of traffic contracts betwwn IPs and infrastructure, develop a solution and conduct experiments.

Integration of the communication architecture with the memory architecture

KTH will develop a scalable, distributed memory architecture for MPSoCs. It will facilitate efficient handling of a virtual address space, cache coherence and memory consistency.

The above milestone has been fulfilled. During the second year Data Management Engine has been developed and implemented, that is programmable and can in principle support all types of MPSoC memory management functions and algorithm.

During year 3 we will develop and implement a programmable memory management handler that realizes many memory management functions efficiently. Empasis will be put on scalable performance for distributed memory system in MPSoCs.

The above milestone has been fulfilled. In the third year, cache coherence protocols, memory consistency models, dynamic memory allocation alrgorithms have been developed.

Focus in year 4 will be on developing a *scalable* cache coherence solution, improve the memory consistency models and their implementations, and conduct more realistic experiments.

Modeling and analysis of heterogeneous systems:

In this new activity, started during year 2, KTH and DTU will develop a SystemC based modeling framework for heterogeneous systems including untimed, synchronous time, discrete time and continuous time models of computation.

The above milestone has been fulfilled partially. The SystemC based models have been created for the synchronous time and discrete time (event driven) MoCs.

KTH and DTU will continue their development of a comprehensive SystemC based modelling framework for heterogeneous systems. Focus in year 4 is on integrating performance analysis tools into the modeling framework, and to implement the framework in SystemC to allow for broader industrial use.

System Level Temperature Modelling and Analysis

During the third year, the Linköping group will work on the elaboration of fast and sufficiently accurate analytical temperature models for the system level. Such an efficient approach to temperature analysis is extremely important as a component in a temperature aware optimisation framework for the design of energy efficient embedded systems.

The above milestone has been fulfilled.

In year 4, LiU will extend the elaborated temperature models to multicore systems.

Simulation-based and analytical methods for performance estimation of distributed real-time systems for control applications

During the third year the groups at Linköping, DTU, and Lund will continue their work on modelling and quality optimisation for control applications implemented on distributed embedded systems. Special emphasis will be placed on the issue of event based control and the related quality vs. resource utilisation tradeoffs.

The above milestone has been fulfilled.

This activity is considered done and will not be continued in year 4.

Modeling and analysis of fault tolerant distributed embedded systems

During the second year Linköping and DTU will develop a reliability analysis approach for distributed and MPSoC systems considering various hardening degrees of the underlying hardware platform.

The above milestone has been fulfilled.

During the third year, DTU will investigate "design for adaptivity". The question is how can a system be designed offline such that runtime adaptation is facilitated. We will identify a relevant case-study that can be used to motivate such an approach and propose design methods for adaptivity.

The above milestone has been fulfilled. The research performed has been reported in the DTU technical report "Fault-Tolerant Design of Mixed-Criticality Adaptive Embedded Systems", by P. K. Saraswat, P. Pop and J. Madsen.

In year 4, DTU will focus on the modelling and analysis of certification costs for mixed-criticality embedded systems. Using such an analysis model, DTU will extend their temporal optimization tool to incorporate trade-offs related to certification costs.

Modeling and Verification of Embedded Systems

In year 1 DTU will continue to formalize the ARTS system-level simulation model using timed automata based on UPPAAL. This work was started in ARTIST2. The aim is to make it usable for designers early in the design process. In order to support designers of industrial applications, the timed-automata model will be hidden for the user, allowing the designer to work directly with the abstract system-level model of embedded systems. The work will be carried out in cooperation with AAU.

The above milestone has been fulfilled, resulting in a prototype framework called MoVES, which allows to experiment with different models-of-computation.

During the second year this work will be continued with the aim to capture more aspects of both the application and the platform. A goal is to make a stronger link between the system-level model and a more detailed hardware platform model. DTU will refine its formal model to address modeling and verification issues closer to the hardware layer of the execution platform.

The above milestone has been fulfilled,

In the third year, DTU will extend the MoVES Framework in various ways, e.g. to incorporate more resource aspects and more advanced bus structures. Furthermore, the issue of scalability will be addressed.

The above milestone has been fulfilled,

In the fourth year, the availability of the MoVES tool will be improved.

Formal verification of design properties of hardware architectures

In year 1 DTU has worked on a formal language for hardware models based on the Gezel hardware description language developed and maintained by Virginia Tech, USA. The aim is to define a clear semantics of the language which allows to formulate the model-of-computation using timed automata, and hence, being able to formally reason about the hardware architecture.

The above milestone has been fulfilled.

In year 2 this work will be continued, with the aim to reason about non-functional properties such as power consumption and memory usage. A number of larger design cases will be carried out.

The above milestone has been fulfilled.

Initial experiments with a refinement framework for Gezel will been conducted. The framework will be elaborated taking resources into account. Furthermore, studies concerning a more semantically streamlined version of Gezel will be conducted.

The above milestone has been partially fulfilled. The studies of a more semantically streamlined version of Gezel has been started but not completed.

This activity will not be continued in year 3.

Analysis tools for embedded systems

In year 1 DTU has established efficient methods for verification of resource constraints. One activity will focus on the real-time logic durations calculus.

The above milestone has been fulfilled. A model-checking result was established which has the potential of verifying strong timing constraints as well as other kinds of resource constraints. The work has been done in collaboration with University of Oldenburg.

Based on the encouraging results from a first prototype implementation, the plan for next year is to extend the theory and to advance the development of the tool.

The above milestone has been fulfilled.

In the third year, DTU will continue the Improvement of the Duration Calculus model checker with respect to both the theoretical aspects and the practical improvements of the prototype.

The above milestone has been fulfilled.

In the fourth year, DTU will continue this activity. Furthermore, investigations concerning implementations exploiting multi-core platforms will be initiated.

Energy Harvesting in Sensor Nodes

ETHZ and Bologna will continue to work on low power sensor nodes. It is planned in year 3 to extend the current application model to different (e.g. distributed) application scenarios. Adaptive clustering and routing will be in the focus of this activity.

This milestone has been achieved. The results are described in the joint publication [MTBB10].

This activity will not be continued in year 4.

Reliability extensions for Variability Aware Modeling (IMEC)

In the 3rd year a research effort will start for the integration of reliability issues like Negative Bias Temperature Instability (NBTI), Hot Carrier Degradation (HCD), Soft Break Down (SBD in oxide) and Soft Errors in the VAM framework.

The above milestone has been fulfilled within the IWT funded Project ELIXIR.

In addition, within the European project SYNAPTIC, regarding lithography evaluation, Imec defined the litho process variability aware methodology for the target technology TSMC40LP to characterize SRAM sensitivity to variability and litho issues.

Within the European project TRAMS, we report a method and its implementation in a prototype tool hereafter called Memory Variability Aware Modeling based on a novel technique that predicts the correct memory wide statistics of any parameter that can be measured in a SPICE/SPECTRE testbench, such as access time, power, stability checks such as read margin, and so on. The method relies on a mix of critical path sensitivities to process variations in its building blocks. Such sensitivity analysis is done at a larger granularity than the transistor level proposed so far for analog circuits, hence leading to a more efficient amount of simulation runs needed hence much less CPU time, and it provides a holistic treatment of all interactions.

In the 4th year, the focus within SYNAPTIC will be on the tool flow and benchmarking while in TRAMS the focus will be on FinFET technology.

Runtime layer design for many-cores architectures (CEA, UNIBO)

During Year 3 CEA LIST has implemented a SW runtime for the management of resources in a multi-core architecture where computing processors are sharing a single memory space with uniform access time properties.

The above milestone has been fulfilled.

In year 4, CEA LIST will tackle the challenge of managing resources of manycores architecture by exending actual framework to support clusterized architectures with a non uniform memory space.

3-D Systems (EPFL)

Since EPFL has only recently joined the cluster, solely future milestones are mentioned

Analysis of multi-clock domain for 3-D systems
Models that include process variability both random and systematic will be developed for different clock distribution networks. Note that the models need to be adapted to the traits of each investigated network.

***IR* drop analysis for 3-D systems**
Effort will be placed to develop a power grid analysis tool that can handle 3-D power didtribution network including the vertical interconnect (a new structure as compared to planar power grids).

3.3 Main Funding

The ArtistDesign NoE funds integration and building excellence with the partners, and with the European research landscape as a whole. Beyond this “glue” for integration and excellence, during Year 3 this activity has benefited from direct funding from:

Linköping University:

- Swedish Foundation for Strategic Research (SSF)
  Project name: “Fault-Tolerant and Secure Automotive Embedded Systems.”
- Swedish research Council

Project name: “Adaptive Resource Allocation for Distributed Embedded Systems.”

DTU

- SYSMODEL (ARTEMIS JU). Period 2009-2011.
- DaNES (Danish Network for Embedded Systems, funded by the Danish Advanced Technology Foundation), Denmark. Period 2007-2010.
- ProCell (project on programmable cell chip: culturing and manipulation of living cells with real-time reaction monitoring funded by the Danish Strategic Research Council),
- RECOMP funded by ARTEMIS JU. Period 20010-2013.
- SMECY funded by ARTEMIS JU. Period 20010-2013.
- ASAM funded by ARTEMIS JU. Period 20010-2013.

CEA LIST

- MC2H (ManyCore for Computing and Healing). French R&D cooperation program (Nano 2012), focusing on the design of multi-core component and on the development of the SW layer allowing to manage it. In this project CEA LIST mainly focus on the SW runtime development in this project.
- SCALOPES (SCAlable LOw Power Embedded platformS). ARTEMIS project. The project focus technology and tool developments for multi-core archictures for communication infrastructure, surveillance systems, smart mobile terminals and stationary video systems.
- SMECY (Smart Multicore Embedded SYstems). ARTEMIS project. The mission of the SMECY project is to develop new programming technologies enabling the exploitation of manycore architectures

UNIBO

- ICT-Project PREDATOR
- ICT-Project GALAXY
- ICT-Project Scalopes (Artemis JU)
- Industrial funding on Sensor Networks from Telecom Italia spa

TUBS

- COMBEST (IST STREP 215543)
  This IST STREP project COMBEST provides a formal framework for component based design of complex embedded systems.
  http://www.combest.eu/home

ETH Zurich

- PROD3D Programming for Future 3D Architecture with Many Cores (EU FP7)
- EURETILE Mapping Algorithms onto Tiled Multirocessor Arrays (EU, FP7)
- PREDATOR Predictable and Efficient Embedded Systems (EU, FP7)
- COMBEST Component Based Design of Embedded Systems (EU FP7)
- MICS Mobile Information and Communication Systems (Swiss National Science Foundation)

KTH

- SYSMODEL (ARTEMIS JU). Period 2009-2011.
- Swedish national funding VR for a NoC performance analysis project. Period 2009-2011.
- iFest (ARTEMIST JU). Period 2010-2013.

EPFL

- PROD3D Programming for Future 3D Architecture with Many Cores (EU FP7)
- Guaranteeing Power and Signal Integrity for 3-D ICs (Swiss National Science Foundation)

- Nanosystems (Advanced ERC grant)

-- Changes wrt Y2 deliverable --

The list of projects funding the activities in the activity, has been updated

4. Internal Reviewers for this Deliverable

- **Dr. Raphaël DAVID** (CEA LIST)
- **Associate Professor Paul Pop** (DTU)