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

Activity Progress Report for Year 3

Platform and MPSoC Design

Cluster:

Hardware Platforms and MPSoC

Activity Leader:

Prof. Luca Benini (UNIBO)

http://www-micrel.deis.unibo.it

Policy Objective (abstract)

The main objective of the activity is to build a common research environment and a wide basis of technical and scientific competences on embedded platforms, with special emphasis on Multi-processor Systems-on-Chip (MPSoCs). The main challenges are to find adequate solutions for the significant fragmentation and lack of integration in this area. The consensus on the fact that hardware platforms for embedded applications will be multi-core, with increasing degrees of parallelism, is not matched at the software and system design level. The teams involved in the activity aim at building stronger common techniques for modeling, analysis and run-time management of embedded hardware platforms. The expected impact will be a faster and more consistently focused development of methods and tools in support of application development and mapping.

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

Team leader: Prof. Lothar Thiele – TIK, ETH Zürich (Switzerland). Role: design methods for MPSoC that combine performance analysis with multi-objective application mapping strategies. To this end, an available programming environment DOL (distributed operation layer) will be enhanced and combined with tools from other partners.

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

Team Leader: Prof. Luca Benini – UNIBO (Italy). Roles: (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.

Team Leader: Prof. Jan Madsen – IMM, Technical University of Denmark (Denmark) Areas of his team's expertise: abstract RTOS and NoC models for multiprocessor system simulation and verification. Modeling and analysis of fault-tolerant embedded systems.

Team Leader: Prof. Rolf Ernst – TU Braunschweig (Germany) Roles: TU Braunschweig contributes methods to deal with variability and reliability issues for systems built from unreliable components.

Team Leader: Prof. Maja D'Hondt – Interuniversity Microelectronics Centre, Imec vzw. (Belgium) Imec contributes to design-time and run-time resource management optimizations for MPSoC platforms.

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-tasks applications onto multi- and many-core devices
(iii) Design of multi-core architectures for dynamic multi-task applications

Team Leader: Prof. Giovanni De Micheli – EPF Lausanne (Switzerland). This team will work on novel models and policies for run-time control of MPSoC platforms. Areas of expertise include hardware design methods and tools, algorithms, real-time systems and 3D integration.

--- Changes wrt Y2 deliverable --

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

1.2 Affiliated participants and their role within the Activity

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.

Prof. Krish Chakrabarty - Duke University (USA).
Areas of his team's expertise: first to develop droplet-based biochips that use electrowetting on dielectric for droplet transport. Design methods and tools for dropletbased biochips.

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. Per Gunnar Kjeldsberg - Norges Teknisk-Naturvitenskapelige Uni., NTNU (Norway)

This team will introduce novel task migration methodologies for MPSoC platforms utilizing hardware accelerators.

Prof. David Atienza – EPFL (Switzerland)

This team will introduce novel run-time memory management optimizations for MPSoC platforms.

Dr. Valter Bella, Telecom Italia Lab (Italy)
Architecture and design of wireless sensor networks and embedded systems for ambient intelligence.

Dr. Daniel Karlsson, Volvo Technology Corporation (Sweden)
Architecture and design of automotive embedded systems.

Dr. Diego Melpignano, ST Microelectronics (Italy)
System SW architecture for many-core consumer platforms

Dr. Eric Flamand, STMicroelectronics (France)
Hardware architecture for many-core consumer platforms

Dr. Matthias Gries, Intel (Germany)
Thermal control and energy management for multi-core platforms

-- Changes wrt Y2 deliverable --

Prof. Dimitrios Soudris from Democritus Uni. of Thrace, DUTH (Greece) has joined as Affiliated Partner

1.3 Starting Date, and Expected Ending Date

Starting date: January 2008.

Ending date: the activity will span the duration of the project, and continue beyond the end of the project. This is because all current trends indicate that MPSoC and platform design will increasingly become a primary concern and focus of action for researchers, designers and developers working on embedded systems.

Moreover, the integration achieved by this activity is creating the know-how and the skills required for the definition of new research and development initiatives. As a result, all partners

are already actively involved in long term funded research programs in MPSoC and embedded design, whose end-date is beyond the duration of ArtistDesign.

-- Changes wrt Y2 deliverable --

No changes with respect to Year 2.

1.4 Policy Objective

While there is wide consensus on the fact that hardware platforms for embedded applications will continue to be multi-core, with increasing degrees of parallelism, the evolution trajectory of programming models, design-time and run-time application environments is much less clear. The consequence is fragmentation: while many research teams are working on one or more of these domains, there is little communication and integration, this leads to duplication of results and overall slow progress. The teams involved in this activity have a wide-ranging research experience which covers all the key areas in MPSoC application specification mapping. The integration activity supported by ARTIST-DESIGN will help the participants to the cluster in strengthening the coherency of their approaches and focus on addressing complementary issues in a synergistic fashion. In particular, there will be an initial effort in reaching a common consensus on the most critical issues to be addressed, define common terminology and decide the operational strategy to address them in a collaborative fashion. The expected impact will be a faster and more consistently focused development of methods and tools in support of application development and mapping.

-- Changes wrt Y2 deliverable --

No changes with respect to Year 2.

1.5 Background

The partners involved in this activity have very active ongoing cooperations on a number of topics. A non-exhaustive set of examples of background cooperation activities is given here.

Imec and UNIBO have ongoing collaboration on dynamic memory management optimizations at the system level for single processor systems, which they plan to extend in the domain of multiprocessor systems.

DTU has strong experience in language constructs that allow the programmer to express communication among tasks in shared memory programming models. MPSoC design and programming has also been performed on hardware structures which makes it possible to expose the memory hierarchy to software and allows multiple programming models to co-exist on the same platform. In addition DTU is collaborating with Duke University (DUKE) on Synthesis for Biochips.

EPFL and UNIBO have ongoing activities in the areas on Networks on Chips (NoC). Specific joint research topics include design methods and tools for 3D networks on chip, with particular interest in wire-planning three-dimensional routes. Another common research interest is the design of NoCs with Quality of Service (QoS) guarantees.

UNIBO and ETH Zurich have ongoing collaborations on Optimization-centric MPSoC Design and on optimal management of smart sensors with energy harvesting capabilities. Wireless sensor networks are a very relevant example hardware platform with very tight energy

constraints. The limited battery lifetime can be extended indefinitely if the node is equipped with energy harvesters that collect and store energy from the environment. However, given the erratic nature of environmental energy sources, the rate at which sensing, computation and storage operations can be performed should be dynamically adjusted to the energy availability using a closed-loop optimal control policy.

Linköping and UNIBO have cooperated on allocation and scheduling policies for low power systems, where clock frequency and voltage setting are also degrees of freedom for optimization.

DTU and Linköping have cooperated on optimisation of distributed embedded systems.

-- Changes wrt Y2 deliverable --

No changes with respect to Year 2.

1.6 Technical Description: Joint Research

The main scientific challenges addressed in this activity are focused on how to map complex applications onto multi-core hardware platforms. This includes addressing allocation and scheduling issues like: scalability, flexibility, composability, predictability, design-time reduction and increased dynamism. The problem is complex and multi-faceted. On one hand, we have static (design/compile time) approaches, where applications are analyzed and optimal mapping decisions are taken before the platform is deployed in the field. On the other hand, we have dynamic, run-time approaches where mapping decisions are taken online, and they are triggered by environmental and workload variations. While these approaches start from different premises, they should not be regarded as alternative, rather they are synergistic.

Design time analysis and decisions can help in providing a good starting point for run-time adaptation, moreover off-line pre-computation can reduce the overhead of the online policies making them more reactive and less resource-hungry. One important requisite for any mapping strategy is to ensure predictability AND efficiency. Note that online adaptation is not adverse to predictability: if online adaptation is based on feedback control (e.g. finite horizon), it can be used to “stabilize” the system, and make it more robust (predictable) in response to environmental variations (e.g. temperature).

Another scientific challenge addressed in this activity is the development of innovative reliable multi-core programming models and architecture platforms able to address computation and control-oriented applications. One key building block is the development of efficient synchronization & communication abstractions that are required for successfully deploying MPSoCs in embedded application domains. Efficiency is inherently related to both power and performance, hence it is an energy metric. In embedded systems, productivity-enhancing abstractions are acceptable only if they do not compromise efficiency, so the focus is on how to enable fast development (debugging, tuning) without losing efficiency. It is also extremely important to take into account variability of both hardware fabrics and application workloads, which is deemed to increase rapidly. Hence, the concurrency management layer should provide means for dynamically managing workload variations, as well as hardware unpredictability sources. Nonetheless, this activity addresses the study of multi-core/many-core systems in relation to 3D integration technologies. We worked on specific hardware abstactions and on run-time policies that perform well on these models.

-- 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.

ETHZ has been mainly involved in cooperations with **UNIBO** and **University Dortmund** on new approaches to map algorithms onto highly parallel MPSoC platforms. To this end, the specification and mapping environment DOL (distributed operation layer) was linked to the MPARM simulation platform from UNIBO.

The **Linköping** group has addressed two major issues: *Design optimisation of fault tolerant distributed real-time systems* and *Energy efficient design of embedded real-time systems*. Linkoping has addressed the problem of energy-efficient design for time constrained multiprocessor systems. In particular, the problem of thermal aware energy optimisation has been tackled. The major challenge was to integrate temperature modelling into the framework of energy efficient system level scheduling and voltage selection.

DTU has addressed issues related to the following macro-topics: (i) *MPSoC design and architectures*. In this context, the emphasis has been on methods to develop MPSoC platforms, covering application specific platforms as well as platforms for dynamic reconfiguration. (ii) *MPSoC programming*. In this context, the emphasis has been on exploring high-level programming models for multi-core architectures and on understanding the dynamic behavior of run-time reconfigurable systems with the aim of developing efficient run-time management algorithms. (iii) *Synthesis for Biochips*. This is a new activity which aims at using principles from MPSoC design to design biochips based on digital microfluidics. Emphasis has been on understanding the biochip platform and identifying the design problems related to it.

UNIBO has addressed, in cooperation with ETHZ, the modelling of miniaturized energy harvesting devices for perpetually powered systems. In particular design of photovoltaic energy harvester for distributed embedded systems was optimized with a methodology which can be easily applied to embedded systems in order to extend battery lifetime. Furthermore it can be used to optimize the design of harvesting ICs; to this end we proposed an inductor-less architecture, suitable for on-chip integration, which permits to increment the conversion process efficiency at the minimum power consumption.

Imec and its affiliated partners (ie, DUTH, UCM and NTNU) have tackled the establishment of a common profiling and run-time MPSoC resource management exploration framework. More specifically, the existing MATISSE and MATADOR frameworks were used as foundation and extended using the software metadata and system scenario approaches. Additionally, in collaboration with DUTH, KTH, TU/e and TU Dortmund a common design flow and tool flow was investigated in order to link the run-time memory optimization methodologies and tools with the design-time memory optimization and source code parallelization methodologies and tools. This work was performed in cooperation with the Software, Synthesis, Code Generation and Timing Analysis cluster and involved additional teams outside the ArtistDesign network.

Some significant achievements obtained by the partners involved in this activity are summarized below.

Temperature Aware System-level Power Optimization (Linköping, UNIBO together)
Linköping and UNIBO have cooperated on system-level design issues, focusing on energy optimization, and more specifically on the optimization of energy-efficient time constrained

multiprocessor systems. In particular, the problem of thermal aware energy optimisation has been tackled.

The major challenge was to integrate temperature modelling into the framework of energy efficient system level scheduling and voltage selection.

High power densities in current SoCs result in both huge energy consumption and increased chip temperature. We have elaborated a temperature-aware dynamic voltage selection technique for energy minimization and performed a thorough analysis of the parameters that influence the potential gains that can be expected from such a technique, compared to a voltage selection approach that ignores temperature. We have also made a study regarding the relevance of taking into consideration transient temperature effects at optimization, the impact of the percentage of leakage power relative to the total power consumed and of the degree to which leakage depends on temperature. Moreover, we have also proposed a temperature-aware task mapping technique for energy optimization in systems with dynamic voltage selection capability.

Run-time resource management (DTU, Imec)
Understanding the dynamic behavior of run-time reconfigurable systems is a very complicated task, due to the often very complicated interplay between the application, the application mapping, and the underlying hardware architecture. However, it is a key issue to determine the right reconfigurable architecture and a matching optimal on-line resource management policy. Although architecture selection, application mapping and run-time system have been studied intensively in the past, they have not been thoroughly studied and modelled in the context of run-time reconfigurable system. DTU has extended its simulation framework COSMOS to study the dynamic behavior of run-time reconfigurable systems. COSMOS is an extension of the ARTS multiprocessor simulation framework, which was developed during ARTIST2.

Through a number of design space exploration experiments, we have pinpointed the critical design issues in the reconfigurable architecture study and analyzed their impact on the architecture performance. Experiments with various run-time resource management polices have shown that it is possible to gain performance from such architectures and have suggested some general guidelines for obtaining efficient run-time resource management.

-- 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)*

Design optimisation of fault tolerant distributed embedded systems (Linköping, DTU)
Linköping University and DTU have developed an approach to the analysis and design of safety critical, fault tolerant embedded applications with soft and hard real-time constraints (for the analysis aspects see “Platform and MPSoC Analysis” activity report). The goal is to guarantee the deadlines for the hard processes even in the case of faults, while maximizing the overall utility. Time/utility functions are used to capture the utility of soft processes. Process re-execution is employed to recover from multiple faults. A single static schedule computed off-line is not fault tolerant and is pessimistic in terms of utility, while a purely online approach, which computes a new schedule every time a process fails or completes, incurs an unacceptable overhead. Thus, a quasi-static scheduling strategy is used, where a set of schedules is synthesized off-line and, at run time, the scheduler will select the right schedule based on the occurrence of faults and the actual execution times of processes. An optimisation

technique for the generation of schedule tables supporting such a quasi-static scheduling approach has been developed and implemented. Moreover, a preemption technique is elaborated as a method to generate flexible schedules that maximize the overall utility for the average case while guaranteeing timing constraints in the worst case. The scheduling algorithm determines off-line when to preempt and when to resurrect processes.

Temperature Aware System-level Power Optimization (Linköping, UNIBO together)
Linköping and UNIBO have cooperated on system-level design issues, focusing on energy optimization, and more specifically on the optimization of energy-efficient time constrained multiprocessor systems. In particular, the problem of thermal aware energy optimisation has been tackled.

The major challenge was to integrate temperature modelling into the framework of energy efficient system level scheduling and voltage selection.

High power densities in current SoCs result in both huge energy consumption and increased chip temperature. We have elaborated a temperature-aware dynamic voltage selection technique for energy minimization and performed a thorough analysis of the parameters that influence the potential gains that can be expected from such a technique, compared to a voltage selection approach that ignores temperature. We have also made a study regarding the relevance of taking into consideration transient temperature effects at optimization, the impact of the percentage of leakage power relative to the total power consumed and of the degree to which leakage depends on temperature. Moreover, we have also proposed a temperature-aware task mapping technique for energy optimization in systems with dynamic voltage selection capability.

Photovoltaic scavenging systems from the model to the optimized design (UNIBO, ETHZ, together)
UNIBO and ETHZ have improved the design of a scavenger prototype which exploits miniaturized photovoltaic modules to perform automatic maximum power point tracking. We propose a detailed model of the solar cell that predicts the instantaneous power collected by the panel and improves the simulation of harvester systems. Furthermore, we focused on a methodology for optimizing the design of MPPT solar harvesters for self-powered embedded systems and presented innovations in the circuit architecture with respect to our previous implementation. We verified that energy consumption and efficiency of the MPP tracker are very important design criteria in energy scavengers for sensor nodes, therefore we analyzed two important metrics: (i) maximization of the energy harvesting efficiency; (ii) minimization of the energy used for ineffective operations.

Modelling and Evaluation of Reliability Analysis for MPSoCs (TU Braunschweig)
TU Braunschweig has taken several steps to address the prediction of performance implications of errors in systems with unreliable components and fault-tolerance mechanisms. The main focus is the impact of these factors on real-time performance. For this purpose initial timing models have been developed to reflect the performance of such systems in the error-free case as well as in case of errors. These models enable an exact specification of the timing effects different fault-tolerance mechanisms may have. In this context we currently restrict our consideration to transient errors, such that systematic design failures can be excluded.

We have developed two methods to evaluate the reliability, i.e. the ability of a system or a component to perform its intended function. For this, we obtain a time-dependant function that describes the probability of correct operation during a period of time. The first approach is Monte-Carlo simulation. Based on a given system configuration the timing behaviour of a

component is simulated, whereas errors are randomly inserted corresponding to the specified error model. Because Monte-Carlo simulation is based on random events a large number of simulation runs has to be performed, what leads to long runtimes until accurate results are obtained.

EPFL

For EPFL, year 2 saw a transition periof from Prof. Henizinger to Prof. De Micheli. Prof. De Micheli was introduced to the Artist2 network and participated to the summer school in Autrans.

No significant activity was performed in 2009 as participation to the Artist Network started only in the fourth quarter. Nevertheless, meetings and exchanges were done between EPFL and UNIBO.

-- 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)*

In 2010, **CEA LIST** has optimized the hardware support for run-time services in symmetric multi-core architectures. By gathering all key features needed from run-time software (those having strong impact on system reactiveness), a low cost hardware component has been designed. In cooperation with **UNIBO**, an optimized implementation of key resource management services, taking benefits from this hardware support, has also been proposed. Also, an asymmetric multi-core device, relying on a specific execution model and based on a flexible control solution has been designed. This asymmetric approach allows for optimizing processing unit occupation rate at the cost of specific execution model. The objective of this dual approach, with symmetric and asymmetric architecture design, is to quantify performance gains coming from a specific execution model towards a more generic approach running on a symmetrical fabric***.*** The work performed in this area is of interest to associate partner STMicroelectronics in the context of their project ***Platform 2012*** for many-core computing in nano-meter technologies.

TU Braunschweig has continued the work on the ***reliability analysis for unreliable components in safety-critical systems*** with its industrial partners (in particular Symtavision and Toyota-ITC). Additionally the academic research project ASTEROID has been started, focusing on new hardware and software concepts to enhance reliability of safety-critical embedded systems based on microkernels.

In 2010, special emphasis has been put on extending the existing analysis methods for the CAN bus to incorporate the system layer. For that purpose two orthogonal problems have been addressed. First, reliability analysis has been adapted to include processing elements, especially multi-cores. Based on the analysis principles developed for the CAN bus, new methods have been developed to take the characteristics of multi-core CPUs into account. The second topic was composition of individual analyses to derive the reliability of the overall system. The existing methods only accounted for reliability analysis of components in isolation, without considering system-wide timing constraints or correlations between the timing behaviour of different components. Assuming a system with two CAN buses connected by a gateway, major challenges of efficient and accurate reliability analysis have been identified and first solutions have been proposed.

ETHZ has been concentrating on the following problems:

- Energy-aware mapping of tasks onto MPSoC platforms (partly joint work with SSSA (Giorgio Buttazzo)) [PCLST10], [SCT10], [CT10], [HSCTB10].
- Design flows for mapping algorithms onto MPSoC platforms (partly jointly with RWTH Aachen (Rainer Leupers)) [BHKT10], [LTNKWI10].
- Multi-objective optimization for energy aware compilation strategies (joint work with University Dortmund (Peter Marwedel) [LPFMT10].

As can be seen from the above results, there has been substantial progress in combining various methods and approaches. Still an open problem is the integration of online adaptivity in the mapping and optimization process.

UNIBO continued to collaborate together with the **Scuola Superiore Sant'Anna** (**SSSA**) to the development of Elastic Scheduling algorithms on a TDMA Bus. The goal is to ensure the highest utilization of the processors even in case of dynamic variations of the workload at run-time. UNIBO achieved a robust system to dynamic workload variations demonstrating the effectiveness of the approach using a very well known Quality of Control Index.

Moreover **UNIBO** addressed also the development of simulation enviroment for GPGPUs and continued the research on Energy Harvesting exploiting the hardware and software capabilities of a multi-source energy harvester.

DTU continued its work on architectures and programming models for MPSoC. During 2010, DTU has participated in the SMECY project and worked on methods and tools that allow programmers to express their domain knowledge. Part of this work has been done in collaboration with IBM Haifa. Focus has been on operating systems and on hardware simulation platforms based on FPGAs.

DTU has developed a prototype of a new reconfigurable hardware platform with self-organizing and self-healing capabilities. Several self-healing strategies have been evaluated, and compared in terms of performance and recovery capabilities.

Linkoping and **DTU** have continued their work on fault tolerant distributed and multiprocessor embedded systems. **Linkoping** has focused on the optimised implementation of error detection techniques. Once such techniques are available, they can be used to quickly detect transient faults. **DTU** has focused on mixed-criticality hard/soft real-time fault-tolerant systems that have to tolerate transient failures in hard tasks, such that hard deadlines are satisfied and the quality of service for the soft tasks is maximized.

Furthermore, the **Linköping** group has continued the research concerning temperature aware and energy efficient design of real-time embedded systems. The emphasis was on elaboration of fast and sufficiently accurate analytical temperature models for the system level. Based on these models efficient online temperature aware leakage energy optimisation techniques were elaborated.

EPFL was very active in Year 3. Major problems that we tackled include, but are not limited to: 1) Network on Chips models and tools; 2) 3D integration models and analysis tools; 3) Study of NoCs for 3D integration. Exchanges with **UNIBO** continued from the previous years. Prof. Benini spent 2 months at **EPFL** as Visiting Professor.

-- 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

Hardware support for run-time management of multi-core architectures (CEA LIST, UNIBO, STMicoelectronics)

CEA LIST has proposed a run-time resource management acceleration solution by means of a **Hardware Synchronizer** (HWS). The HWS includes a synchronization module, an event and an interrupt generator, a dynamic allocator, and a fault-tolerance module. The synchronization module provides a hardware-supported acceleration of various synchronization mechanisms: semaphores, mutexes, barriers, joins, etc. The event and interrupt generators allow any processor to generate an event or interrupt to any other processing element (PE) in a shared memory architecture. To avoid systematic polling operations, automatic notification is made possible via the programmable notifier which can send trigger events on specific synchronization values. The interrupt generator also contains a system messenger providing a set of mail boxes to generate interrupt notifications to a target PE, whenever a new message is stored. The Dynamic Allocator allows the system to dynamically assign a task to an available PE. It is able to check the resource states (PE fault status, dynamic interrupt priority level table, etc.) and to allocate the best processor to execute the task. The fault-tolerance module gives access to the faulty states of cluster cores and controls a set of watchdog timers for detecting violations of execution deadlines in the application due to transient or permanent fault occurrences in the cluster.

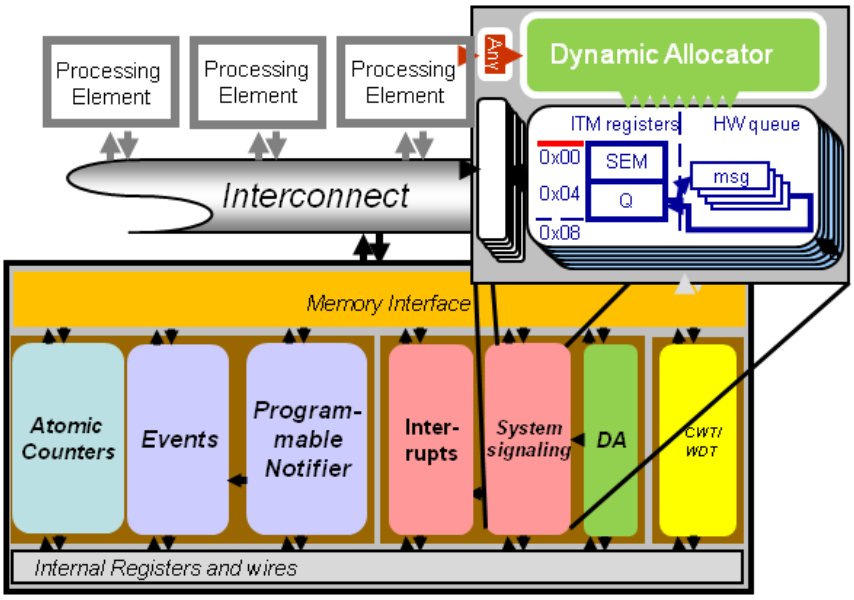

To minimize area overhead of a computing cluster, the HWS has a low silicon footprint with a complexity limited to approximately 120kgates, which represents less than 1% of the overall area of a computing cluster made of 16 processors and a 16 shared memory. On the basis of this module, optimized resources management services can be implemented. For example, simple reactive task management has been implemented in cooperation with UNIBO and has proven to allocate tasks to computing resources in tens of cycles. For a VC1 video decoding application, parallelized at a very fine grain, performance overhead has been shown to be less than 3%.

CEA LIST also works on the FPGA prototype of its SCMP asymmetric multi-core platform. Especially CEA LIST has designed a programmable controller for an asymmetric multi-core

architecture based on the small RISC processor (AntX) coupled to a specific accelerator to accelerate control operations like sorting. Also a specific memory management unit, which can support a dynamic management of buffers for concurrent readers and writers to the same buffer, has been prototyped. A small OS kernel running on it to manage the execution of dynamic dataflow applications has been implemented to take benefits from this hardware support. This activity is of interest to associate industrial partner STmicroelectronics in the context of the Platform 2012 project

Design of High-Efficiency Energy Harvester for Autonomous Embedded Systems (UNIBO)

UNIBO continued the research on Energy Harvesting to enable systems operating in a deploy and forget mode. After Photovoltaic scavenging systems addressed in Y1 and Y2, UNIBO moved to optimize the design of micro turbine energy harvesters which exploit air flow as energy source. A detailed model has been defined, and the final prototype outperfoms the state of the art. The core of the harvester is a high efficiency buck-boost converter which performs an optimal power point tracking. UNIBO expressly designed and optimized the DC-DC converter and the Maximum Power Point Tracking (MPPT) circuit, to provide the resistive load which maximizes its performance. The measured efficiency of the converter is always around 87%. The adopted design methodology enhances the solution proposed in the literature and is aimed at the minimization of the power losses of the operating devices. This procedure allows to make the best design choices on the basis of the expected operating conditions of the harvester.

Adaptive TDMA bus Allocation and Elastic Scheduling (UNIBO, SSSA)

Predictability and modularity of MPSoC for Real-Time applications are open issues in the research community. The main challenges are represented by the contention due to shared resources, such as a shared bus, and by the inner complexity of the whole system, both from the architectural and SW application point of view. Several techniques were introduced to address these issues, such as TDMA scheduling schemas for resource sharing, or the adoption of Elastic Scheduling techniques to relax some bounds on the period of the task-set. Anyhow, they further increase system complexity at each level, and result in an overall worsening of the WCET (Worst-Case Execution Analysis) of running tasks, producing conservative results.

UNIBO together with **SSSA**, designed a system where Elastic Scheduling and the TDMA Bus work synergistically to ensure the highest utilization of the processors even in case of dynamic variations of the workload at run-time. When a processor is subject to a workload change, it makes a request to a Master Processor to adapt its share of bus bandwidth acting on the TDMA scheduling. The Master mediates among requests and produces a new allocation of the shared resource. As a consequence of the new bus scheduling, memory latencies change for all the processors, and so do task execution times as well. Thus, WCET analysis must be recomputed in a fast and light way and task periods may be adjusted, by the adoption of Elastic Scheduling, to reach a desired (typically maximum) utilization of the processors.

UNIBO implemented the technique on an MPSoC cycle accurate simulator and performed extensive experiments to validate the proposed approach and to test its effectiveness with real case automotive and avionics application tasks. The resulting system is robust to dynamic workload variations, and we showed the effectiveness of our approach using a very well known Quality of Control Index.

Instruction set simulator for thousand-core architectures running on GPGPUs (UNIBO, EPFL)

Future architectures will expose thousands of simple processing cores and on-chip memories connected through network on chip, which speed is hundreds times faster than the off-chip one. **UNIBO** and **EPFL** ESL lab developed a novel parallel simulation technique which exploits the flexibility of modern GPGPUs. The simulator is capable of modeling thousands of computational units with private data and instruction caches, connected through a network on chip. At the state of the art we support X86 and ARM ISSs, to use the simulator to both HPC (High Performance Computing) and embedded domains. The Cache Simulator can model different cache configurations with different associativity and writing policies. Finally the network simulator connects all computational nodes within a mesh, allowing communication through shared memory.

The modularity of our implementation enables the exploration of different target architectures, starting from very simple nodes with data and instructions SPM (ScratchPad Memories) up to complex computational nodes equipped with both data and instruction caches. Experimental results show that the system achieves good performance, ~1000 MIPS (Millions Instructions per Second) simulating 1024 target cores.

Policies for thermal control (EPFL, UNIBO, Intel Braunschweig)

EPFL studied model predictive control methods for designing run-time policies for controlling temperatures in systems on chips. Objectives are: 1) reducing hot spots; 2) reducing spatial and temporal gradients of temperature; 3) keeping the chip temperature within a given band; 4) satisfying performance requirements, were performance is measured in terms of done/undone work within a given time window. Four methods were developed, tested and compared. Data analysis showed respective merits and disadvantages of these policies. Other activities addressed quality of service for networks on chips. In particular, we studied bounds for performance that can be inferred from the structure of a network. The activities on thermal control are of interest to Intel for their high-performance multi-core chips.

Heterogeneous Multi-core Platform for Software-Defined Radio (Imec vzw)

A software-defined radio (SDR) system is a radio communication system in which physical layer components are implemented on a programmable or reconfigurable platform. The modulation and demodulation is performed in software and thus the radio is able to support a broad range of frequencies and functions concurrently. Imec's SDR platform is capable of running various wireless standards like: Wireless LAN, LTE, WiMAX etc. However, in order to reach an energy efficient yet programmable solution for different radio standards, the designer is required to make the trade off between flexibility and programmability. The solution presented in this work is a heterogeneous multi-core platform.

An RDL-Configurable 3D Memory Tier to Replace On-Chip SRAM (Imec vzw)

In conventional SoC designs, on-chip memories occupy more than the 50% of the total die area. 3D technology enables the distribution of logic and memories on separate stacked dies (tiers). This allows redesigning the memory tier as a configurable product to be used in multiple system designs. Imec proposes a one-time configurable memory tier designed to minimize the performances overhead due to the commodity. Flexible configuration is enabled by smart memory macros and I/Os organization and a customizable redistribution layer routing. With respect to the dynamic reconfigurability, the proposed design offers up to 40% faster access time, while saving more than 10% of energy per access. In addition, production cost trade offs are analyzed.

Reliabiltiy Analysis with mixed-criticality workloads (TU Braunschweig, Toyota-ITC, Symtavision) *(see also report section Integration Driven by Industrial Applications)*

Considering reliability analysis for multi-core CPUs, promising results have been obtained in year 3. Based on a generic multi-core platform that applies checkpointing and rollback mechanisms to protect safety-critical processes against faults, an accurate analysis algorithm with quite good performance has been developed. Especially the issue of mixed-criticality design has been addressed. Results are going to be published soon.

Even though still under development, first results exist in the context of system-wide reliability analysis. Two different solutions have been proposed to derive the reliability of a gated CAN-to-CAN network, with considerable differences in computational complexity and accuracy. The first solution follows the paradigm of strict composability. System timing constraints are broken down to component layer such that the component-based approaches can be reused. Results are combined afterwards to obtain the overall system reliability. This approach has very good performance. However, it is not totally accurate due to the abstraction of important timing relations at component interfaces. Further on, it is very restrictive with respect to the considered gateway architecture. The second approach is built on a holistic view of the system, without abstraction of component interfaces. However, it is expected to have worse performance compared to the compositional one.

Furthermore the research on CAN bus reliability analysis has been intensified in year 3. In this context the effect of local clock drift in distributed embedded system has been considered. Results are submitted for publication.

Architectures for on-chip communication in future multi- and many-core processors (TU Braunscheig, Intel Labs)

In the COMPOSE Project, the TU Braunschweig cooperated with Intel Braunschweig on new architectures for on-chip communication in future multi- and many-core processors. The goal was to develop predictable communication mechanisms and service guarantees for real-time and streaming applications. In existing embedded MPSoC, predictable communication comes at the cost of increased latencies for regular best-effort traffic, because this traffic class is treated as a "second class citizen". At the same time, a prioritization of traffic with guaranteed throughput requirements is not beneficial, because streaming applications are usually very latency-tolerant due to their predictable access patterns. TU Braunschweig has proposed and developed architectures for efficient combination of best-effort and real-time traffic. The key idea is to prioritize best-effort traffic for optimal latency, but limit its rate to retain throughput guarantees of real-time traffic. For the rate limitation, two alternatives have been developed. The first is based on distributed traffic shaping [DEK10], while the second exploits the buffer occupancy of real-time traffic to limit best-effort traffic. These mechanisms result in latency improvements of up to 47% compared to traditional real-time capable architectures. The throughput guarantees have been formally proven using a compositional performance analysis approach similar to the SymTA/S approach developed at TU Braunschweig.

The COMPOSE project has been successfully concluded in mid-2010. Some of its results have been presented and discussed at the Artist Design Workshop “Mapping Applications to MPSoCs 2010” and at the Intel European Research and Innovation Conference. *(see also Section 2.5 Keynotes, Workshops, Tutorials)*

The architecture developed in the COMPOSE project also forms the basis of a many-core research platform which is currently developed by TU Braunschweig in the RECOMP project. This platform will be used to implement and evaluate mechanisms (both hardware and software) that support the simultaneous execution of applications with different safety criticalities.

Energy-aware mapping of tasks (ETHZ and SSSA)
The results of these activities have been described in [PCLST10], [SCT10], [CT10], [HSCTB10]. Multiprocessor SOC platforms have been adopted for a wide range of high performance applications, like automotive and avionic systems. Task assignment and processing unit allocation are key steps in the design of predictable and efficient embedded systems. Given the execution modes of applications, we propose a methodology to compute a task to processing element mapping, such that the expected average power consumption is minimized. Changing usage scenarios are represented by varying execution probabilities of modes. Statically pre-computed template mappings for each execution probability are stored on the system and applied at run-time, allowing the system to adapt to changing environmental conditions.

Hierarchical design flow (ETHZ, RWTH Aachen, UBS)
The results of this activity have been published in [BHKT10], [LTNKWI10], [PRTLR10]. Wireless multimedia terminals are among the key drivers for MPSoC platform evolution. Heterogeneous multiprocessor architectures achieve high performance and can lead to a significant reduction in energy consumption for this class of applications. However, just designing energy efficient hardware is not enough. Programming models and tools for efficient MPSoC programming are equally important to ensure optimal platform utilization. Unfortunately, this discipline is still in its infancy, which endangers the return on investment for MPSoC architecture designs. On one hand, there is a need for maintaining and gradually porting a large amount of legacy code to MPSoCs. On the other hand, special C language extensions for parallel programming, as well as adapted process network programming models, provide a great opportunity to completely rethink the traditional sequential programming paradigm for sake of higher efficiency and productivity. MPSoC programming is more than just code parallelisation, though. Besides energy efficiency, limited and specialized processing resources, and real-time constraints also growing software complexity and mapping of simultaneous applications need to be taken into account. In the task, we analyzed the programming methodology requirements for heterogeneous MPSoC platforms and introduced new approaches such as hierarchical representations of event streams.

Multi-objective optimization in Compilation (ETHZ and University Dortmund)
The results of this activity have been published in [LPFMT10]. With the growing complexity of embedded systems software, high code quality can only be achieved using a compiler. Sophisticated compilers provide a vast spectrum of various optimizations to improve code aggressively w. r. t. different objective functions, e. g., average-case execution time (ACET) or code size. Due to the complex interactions between the optimizations, the choice for a promising sequence of code transformations is not trivial. Compiler developers address this problem by proposing standard optimization levels, e. g., O3 or Os. However, previous studies have shown that these standard levels often miss optimization potential or might even result in performance degradation. In this paper, we propose the first adaptive WCET-aware compiler framework for an automatic search of compiler optimization sequences which yield highly optimized code. Besides the objective functions ACET and code size, we consider the worst-case execution time (WCET), which is a crucial parameter for real-time systems. To find suitable trade-offs between these objectives, stochastic evolutionary multi-objective algorithms identifying Pareto optimal solutions are exploited. A comparison based on statistical performance assessments is performed which helps to determine the most suitable multi-objective optimizer. The effectiveness of our approach is demonstrated on real-life benchmarks showing that standard optimization levels can be significantly outperformed.

Energy efficient embedded system design (Linköping)

Technology scaling and ever increasing demand for performance have resulted in high power densities in current circuits, which have also led to increased chip temperature. At the same time, leakage energy has become the dominant energy consumption source of circuits. Due to the strong dependence of leakage current on temperature, growing temperature leads to an increase in leakage current and, consequently, energy, which, again, produces higher temperature. At system level, dynamic voltage selection (DVS) is one of the preferred approaches for reducing the overall energy consumption. This technique exploits the available slack time to achieve energy efficiency by reducing the supply voltage and frequency such that the execution of tasks is stretched within their deadline or by switching the processor to the idle state. During the available slack interval, the processor remains idle and can be switched to a low power state. Due to the strong inter-dependence between leakage power and temperature, different distributions of idle time will lead to different temperature distributions and, consequentially, energy consumption. In this work we have addressed the issue of optimizing leakage energy consumption through distribution of the slack time.

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. In order to support our temperature aware energy optimisation techniques we have elaborated a fast and accurate temperature analysis technique, can be used inside a temperature aware system level optimization loop (for the analysis and modelling aspects see “Platform and MPSoC Analysis” activity report).

Design of fault tolerant distributed embedded systems (Linköping, DTU)

a) Optimisation of Fault-tolerant Systems With Hard and Soft Real-time Constraints

We have elaborated a technique for scheduling of fault-tolerant applications composed of soft and hard real-time processes running on distributed systems. We use process re-execution to tolerate transient faults. We propose a quasi-static scheduling algorithm that generates off-line a tree of fault-tolerant schedules that maximize the quality of service of the application and, at the same time, guarantee deadlines for hard processes. At run time, the online scheduler, with very low online overhead, would select the appropriate schedule based on the occurrence of faults and the actual execution times of processes. The proposed approach can be useful on any distributed system whose worst-case communication delays can be obtained. Examples of such systems include multimedia systems in mobile phones, media players, TV-sets, and automotive infotainment. Our approach is also applicable for safety-critical systems used, for example, in factory automation or automobiles.

b) Optimised Implementation of Error Detection Techniques

Safety-critical applications must function correctly even in the presence of faults. Factors like high complexity, smaller transistor sizes, higher operational frequencies and lower voltage levels have contributed to the increase in the rate of transient and intermittent faults in modern electronic systems. Error detection is crucial for meeting the required reliability of the system. Unfortunately, it is also a major source of time overhead. In order to reduce this overhead, one possible approach is to implement the error detection mechanisms in hardware, which, however, increases the overall cost of the system. Because error detection incurs high overheads, optimizing it early in the design phase of a system can result in a big gain, and often makes the difference between a feasible solution and an infeasible one. We have developed a technique for the optimization of error detection implementation in the context of fault-tolerant embedded systems. We have proposed two optimization algorithms: one considering that, when implemented in hardware, error detection is deployed on static

reconfigurable FPGAs and the other one assuming 1-dimensional (1D) or 2-dimensional (2D) partial dynamic reconfiguration capabilities of the FPGAs.

MPSoC design and programming (DTU)

Current programming languages do not allow programmers to efficiently describe their knowledge when they write parallel programs. This leads to sub-optimal performance and excessive programming effort. DTU has worked on methods and tools that allow programmers to express their knowledge [LKM10, LKM]. A new programming tool for optimizing the performance of parallel programs has been developed in collaboration with IBM in Haifa [LLKZ10, LLKZ11]. The tool extracts information from the GCC compiler and presents it to the programmer via the Eclipse IDE. The tool also, based on the information from the compiler, identifies source code constructs which limits the ability of the compiler to optimize executable code. The tool furthermore provides alternatives for these constructs and can optionally automatically apply changes to the source code.

Existing operating systems do not scale to multi-core architectures. One important reason for this is that the algorithms and data structures used do not allow for efficient parallel execution. DTU has worked on data structures that allow the algorithms in operating systems to scale better. Initial work on scalable priority queues for operating systems has been done [PK11]. The work has provided significant insights into fundamental overheads and performance trade-offs in modern multi-core architectures.

The current state-of-the-art for designing multi-core architectures is to use software simulators. These simulators are very compute intensive leading to long simulation times, especially for large simulated systems. DTU has, during 2010, developed a new high-performance FPGA CPU core, Tinuso [SK10].

Synthesis for Biochips (DTU together with DUKE)

Existing DMB synthesis approaches consider that on-chip operations, such as splitting a droplet of liquid, are perfect. However, these operations have variability margins, which can impact the correctness of the biochemical application. DTU has extended existing design methods from DUKE university to consider that a split operation, which goes beyond specified variability bounds, is faulty. The fault is detected using on-chip volume sensors. DTU has proposed an abstract model for a biochemical application, consisting of a sequencing graph, which can capture all the fault scenarios in the application. Starting from this model, DTU has proposed a synthesis approach that, for a given chip area and number of sensors can derive a fault-tolerant implementation. This research [AMPM10] shows that, by taking into account fault-occurrence information, better quality implementations can be derived, which leads to shorter application completion times, even in the case of faults.

In DMBs, during the execution of an operation, a virtual device can be reconfigured to occupy a different group of cells on the array, forming any shape, not necessarily rectangular. DTU has proposed a Tabu Search metaheuristic for the synthesis of digital microfluidic biochips [MPM10], which, starting from a biochemical application and a given biochip architecture, determines the allocation, resource binding, scheduling and placement of the operations in the application. In our approach, we consider changing the device to which an operation is bound during its execution, to improve the completion time of the biochemical application. The completion time of non-rectangular devices is based on an analytical method that takes into account the amount of operation execution performed during droplet routing.

Regarding Continous Fluidic Biochips (CFBs), DTU has proposed an approach to the system-level modeling and simulation of a cell culture microfluidic biochip [MPMHD10] called ProCell, Programmable Cell Culture Chip. ProCell contains a cell culture chamber, which is envisioned

to run 256 simultaneous experiments (viewed as a 16 x 16 matrix). DTU has proposed a biochip architecture model and a comprehensive fault model that captures permanent faults occurring during chip operation. Using the proposed modeling and simulation framework, DTU has performed an architectural level evaluation of two cell culture chamber implementations. A qualitative success metric has been proposed to evaluate chip performance in the presence of partial failures. The results show that significant improvements in efficiency can be obtained using redundancy, providing improved chances to complete an experiment even in the presence of faults. This decreases the experiment repetition rate while increasing system productivity, saving time and reducing costs.

Run-time resource management (DTU)

DTU has continued its research on adaptivity and autonomy in embedded systems. Focus has been on developing a prototype of the eDNA platform. The prototype has been completed and integrated with the modelling and simulation framework developed in Y2. The prototype platform now fully supports both self-organisation and self-healing. Experiments with up to 81 cores have successfully been tested. A research collaboration with NASAs Jet Propulsion Laboratory in Pasadena (CA) was initiated. The aim of the collaboration was to integrate and implement the control and data processing algorithms of a Liquid Crystal Waveguide Fourier Transform Spectrometer on the eDNA prototype platform. This application is used to detect various gasses, such as methane, and thus is an important instrument in the search for bacterial life, for instance on Mars. The FTS application contains both signal processing and control computation that makes it a very suitable test case for the eDNA platform. A prototype was programmed on a Xilinx Virtex 5 FPGA and was later ported to a National Instruments CompactRIO embedded platform in order to interact with the spectrometer of NASA. This work has resulted in three publications [BMK10, BMK11, BKMLC11]

The work on run-time resource management has also been reported in the transversal activity on design for adaptivity.

Component-based service model (DTU together with B&O ICEpower)

The research on component-based service models focuses on the development of a modelling framework for system level performance estimation of embedded systems including Multi-Processor System on Chip (MPSoC) based configurations. The goal is to provide a framework which supports models described at multiple levels of abstraction which will allow designers to perform design space exploration at the various design stages ranging from initial high level specifications to detailed, bit true, cycle accurate models.

During year 3, DTU and B&O ICEpower have primarily focused on performing a number of modifications and adding extensions to the basic service model as well as implementing these in the performance estimation framework in software. These extensions and changes were developed based on the elaborate case study that was performed throughout year 2 in which the mobile audio signal processing platform from ICEpower was investigated. Since the case study was performed, the mobile audio signal processing platform has been implemented in a product at ICEpower and thus it is now possible to extract actual performance metrics for the purpose of comparison. As a test case for the framework, a simplified version of an ARM7 processor has been successfully modeled.

As an impact of this work, the modelling framework for system level performance estimation of embedded systems is going to be the core system level model in a new ARTEMIS JU project on Automatic Architecture Synthesis and Application Mapping (ASAM), where DTU will collaborate with STMicroelectronics, SiliconHive, ACE compilers, and TUe among others.

The work on component-based service modeling has also been reported in the transversal activity on industrial applications.

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

2.2 Individual Publications Resulting from these Achievements

CEA LIST

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

N. Ventroux, T. Sassolas, R. David, G. Blanc, C. Bechara, SESAM Extension For Fast MPSoC Architectural Exploration And Dynamic Streaming Applications. 18th IEEE/IFIP International Conference on VLSI and System-on-Chip (VLSI-SoC), 2010, Madrid, Spain

TU Braunschweig

[DEK10] Jonas Diemer, Rolf Ernst, and Michael Kauschke, "Efficient Throughput-Guarantees for Latency-Sensitive Networks-On-Chip," in Proceedings of the Asia and South Pacific Design Automation Conference (ASP-DAC 2010), January 2010

[DE10a] Jonas Diemer and Rolf Ernst, "Back Suction: Service Guarantees for Latency-Sensitive On-Chip Networks," in Proceedings of the 4th ACM/IEEE International Symposium on Networks-on-Chip (NOCS'10), May 2010

ETHZ

[PCLST10] Simon Perathoner, Jian-Jia Chen, Kai Lampka, Nikolay Stoimenov, Lothar Thiele: Combining Optimistic and Pessimistic DVS Scheduling: An Adaptive Scheme and Analysis. IEEE/ACM International Conference on Computer-Aided Design (ICCAD), IEEE, San Jose, California, USA, November, 2010.

[SCT10] Andreas Schranzhofer, Jian-Jia Chen, Lothar Thiele: Dynamic Power-Aware Mapping of Applications onto Heterogeneous MPSoC Platforms. IEEE Transactions on Industrial Informatics, IEEE, Vol. 6, No. 4, pages 692 -707, November, 2010.

[CT10] Jian-Jia Chen, Lothar Thiele: Energy-Efficient Scheduling on Homogeneous Multiprocessor Platforms. 25th ACM Symposium on Applied Computing , March, 2010.

[BHKT10] Iuliana Bacivarov, Wolfgang Haid, Kai Huang, Lothar Thiele: Methods and Tools for Mapping Process Networks onto Multi-Processor Systems-On-Chip. Handbook of Signal Processing Systems, Springer, pages 1007-1040, October, 2010.

UNIBO

Davide Carli, Davide Brunelli, Davide Bertozzi and Luca Benini, A High-efficiency Wind-flow Energy Harvester Using Micro Turbine, in: 20th International Symposium on Power Electronics, Electrical Drives, Automation and Motion, Pisa, Italy, 2010

Andrea Bartolini, Matteo Cacciari, Andrea Tilli, Luca Benini, Matthias Gries; **A virtual platform environment for exploring power, thermal and reliability management control strategies in high-performance multicores**, GLSVLSI '10 Proceedings of the 20th symposium on Great lakes symposium on VLSI 2010, May 16 -18-2010, Providence, Rhode Island (USA), on page(s): 311 – 316, ISBN: 978-1-4503-0012-4

EPFL

F. Zanini, D. Atienza Alonso, G. De Micheli, and S. P. Boyd, ``Online Convex Optimization-Based Algorithm for Thermal Management of MPSoCs,'' in *Proceedings of the 20th ACM Great Lakes Symposium on VLSI (GLSVLSI 2010)*, vol. 1, (New York), pp. 203-208, ACM Press, 2010.

F. Zanini, D. Atienza, C. N. Jones, and G. De Micheli, ``Temperature Sensor Placement in Thermal Management Systems for MPSoCs,'' in *Proceedings of the IEEE International Symposium on Circuits and Systems (ISCAS 2010)*, vol. 1, pp. 1065-1068, IEEE Press, 2010.

F.Zanini, C. N. Jones, D. Atienza, and G. De Micheli, ``Multicore thermal management using approximate explicit Model Predictive Control,'' in *Proceedings of the of IEEE International Symposium on Circuits and Systems (ISCAS 2010)*, vol. 1, pp. 3321-3324, IEEE Press, 2010.

Linköping

[LEPI2010] A. Lifa, P. Eles, Z. Peng, V. Izosimov, Hardware/Software Optimization of Error Detection Implementation for Real-Time Embedded Systems, Proceedings of International Conference on Hardware/Software Codesign and System Synthesis (CODES+ISSS 2010), Scottsdale, AZ, USA, October 24-29, 2010.

[BAEP2010] M. Bao, A. Andrei, P. Eles, Z. 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.

[TBEP2010] B. Tanasa, U. Bordoloi, P. Eles, Z. Peng, Scheduling for Fault-Tolerant Communication on the Static Segment of FlexRay, Proceedings of 31st IEEE Real-Time Systems Symposium (RTSS10), San Diego, CA, USA, November 30-December 3, 2010.

DTU

[LKM10] P. Larsen, S. Karlsson, J. Madsen, "Expressing Inter-task Dependencies between Parallel Stencil Operations", in Proceedings of MULTIPROG'10, 2010

[SK10] P. Schleuniger, S. Karlsson, "Tinuso: A processor architecture for a multi-core hardware simulation platform", in Proceedings of MCC'10, 2010

[LLKZ10] P. Larsen, R. Ladelsky, S. Karlsson, A. Zaks, "Compiler Driven Code Comments and Refactoring", in Proceedings of MCC'10, 2010

[LKM] P. Larsen, S. Karlsson, J. Madsen, "Expressing Coarse-grain Dependencies among Tasks in Shared Memory Programs", Submitted to IEEE Transactions on Industrial Informatics

[LLKZ11] P. Larsen, R. Ladelsky, S. Karlsson, A. Zaks, "Compiler Driven Code Comments and Refactoring", to appear in Proceedings of MULTIPROG'11, 2011

[PK11] S. Passas, S. Karlsson, "Comparing the Overhead of Lock-based and Lock-free Implementations of Priority Queues", to appear in Proceedings of MULTIPROG'11, 2011

[SPM10] Prabhat Kumar Saraswat, Paul Pop, Jan Madsen, "Task Mapping and Bandwidth Reservation for Mixed Hard/Soft Fault-Tolerant Embedded Systems," RTAS, pp.89-98, 2010 16th IEEE Real-Time and Embedded Technology and Applications Symposium, 2010

[AMPM10] Mirela Alistar, Elena Maftei, Paul Pop, Jan Madsen, "Synthesis of biochemical applications on digital microfluidic biochips with operation variability", in Proceedings of the 2010 Symposium on Design Test Integration and Packaging of MEMS/MOEMS. pp. 350-357, 2010

[MPM10] Elena Maftei, Paul Pop, Jan Madsen, "Tabu Search-Based Synthesis of Digital Microfluidic Biochips with Dynamically Reconfigurable Non-Rectangular Devices", in Journal of Design Automation for Embedded Systems, Volume 14, Number 3, pp. 287-307, 2010

[MPMHD10] Minhass, Wajid Hassan; Pop, Paul; Madsen, Jan; Hemmingsen, Mette; Dufva, Martin System-level modeling and simulation of the cell culture microfluidic biochip ProCell Presented at: 2010 Symposium on Design Test Integration and Packaging of MEMS/MOEMS, 2010 In: 2010 Symposium on Design Test Integration and Packaging of MEMS/MOEMS : IEEE, 2010

[BMK10] Boesen, Michael R., Madsen, Jan and Keymeulen, Didier. *Integration of the Self-Healing eDNA Architecture in an Embedded System and Evaluation of it Using a Fourier Transform Spectrometer Instrument Application.* Presentation at the ReSpace/MAPLD 2010 conference.

[BMK11] Boesen, Michael R., Madsen, Jan and Keymeulen, Didier. *Autonomous Distributed Self-Organizing and Self-Healing Hardware Architecture – the eDNA concept.* Accepted at the IEEE Aerospace Conference 2011 to be held in Big Sky, MT in March 2011.

[BKMLC11] Boesen, Michael R., Keymeulen, Didier., Madsen, Jan., Lu, Thomas and Chao, Tien-Hsin. *Integration of the Reconfigurable Self-Healing eDNA Architecture in an Embedded System.* Accepted at the IEEE Aerospace Conference 2011 to be held in Big Sky, MT in March 2011.

[BM10] Boesen, Michael R., and Madsen, Jan, *Patent Application (WO2010060923): Biologically Inspired Hardware Cell Architecture*

[THM11] Anders Tranberg-Hansen and Jan Madsen, *A Service Based Simulation Framework for Performance Estimation of Embedded Systems,* To appear in CRC Press Book on "Real-time Simulation Technologies: Principles, Methodologies, and Applications" by Katalin Popovici and Pieter Mosterman. 2011

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

2.3 Interaction and Building Excellence between Partners

Interaction between CEA LIST and UNIBO in this activity is based on face-to-face meetings in Saclay and Genoble and numerous phone meetings to clarify technical aspects regarding symmetric architecture design or run-time software.

Concerning the interaction between **EPFL** and **UNIBO**, Prof. Benini from UNIBO has been a visiting professor at EPFL. Moreover, the technical achievements are based on a set of joint meetings and visits between the research groups.

The interaction in TDMA bus Allocation and Elastic Scheduling activity has been between **UNIBO** and **SSSA** as follows:

- Several phone meetings/discussions between UNIBO and SSSA during 2010.
- The Erika RTOS from SSSA has been extensively used on the MPARM platform developed by UNIBO. This has led to frequent interactions between the groups.

There has been a close cooperation between **ETHZ** and ArtistDesign Partners RWTH Aachen (Rainer Leupers), **TUBS** (Rolf Ernst), SSSA (Giorgio Buttazzo) as well as University Dortmund (Peter Marwedel) concerning various aspects of MPSoC Design. The results of these activities have been published in [PRTLR10], [LTNKWI10], [HSCTB10] and [LPFMT10]. These joint activities concern methods to map applications onto MPSoC platforms and analyze the corresponding properties such as timing and power. In addition, the activities span several levels of abstraction starting from compilation via programming models towards high-level performance models.

-- Changes wrt Y2 deliverable --

Some partners (e.g. ETHZ) have extended their interaction with other partners.

2.4 Joint Publications Resulting from these Achievements

CEA LIST, UNIBO.

M. Ojail, K. Ben Chehida, Y. Lhuillier, R. David, L. Benini. Synchronous Reactive Task Management for Fine Grain Parallelism in Many-cores Architectures. Submitted to Workshop on Parallel Programming and Run-time Management Techniques for Many-core Architectures 2010, Hannover

UNIBO – EPFL

Shivani Raghav, Martino Ruggiero, David Atienza, Christian Pinto, Andrea Marongiu and Luca Benini, “Scalable instruction set simulator for thousand-core architectures running on GPGPUs”, in: High Performance Computing and Simulation (HPCS), 2010 International Conference on, pages 459 -466, 2010

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, 2010.]

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, 2010.

J. Joven Murillo, A. Marongiu, F. Angiolini, L. Benini, and G. De Micheli, ``Exploring Programming Model-driven QoS Support for NoC-based Platforms,'' in *Proceedings of*

International Conference on Hardware/Software Codesign and System Synthesis (CODES+ ISSS'10), pp. 65-74, 2010.

G. De Micheli, C. Seiculescu, S. Murali, L. Benini, F. Angiolini, and A. Pullini, ``Networks on Chips: from Research to Products,'' in *Proceedings of the 47th Design Automation Conference (DAC 2010)*, vol. 1, pp. 300-305, 2010.

C. Seiculescu, S. Murali, L. Benini, and G. De Micheli, ``A Method to Remove Deadlocks in Networks-on-Chips with Wormhole Flow Control,'' in *DATE 2010*, 2010.

UNIBO - ETHZ

Clemens Moser, Lothar Thiele, Davide Brunelli and Luca Benini, Adaptive Power Management for Environmentally Powered Systems (2010), in: Computers, IEEE Transactions on, 59:4(478 -491)

UNIBO -SSSA

P. Burgio, M. Ruggiero, F. Esposito, M. Marinoni, G. Buttazzo and L. Benini, “Adaptive TDMA bus Allocation and Elastic Scheduling: a unified approach for enhancing robustness in multicore RT systems”, 28th International Conference of Computer Design (ICCD), Amsterdam, 2010.

ETHZ-TUBS-SSSA-Aachen

[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.

[LTNKWI10] Leupers, R.; Thiele, L.; Xiaoning Nie; Kienhuis, B.; Weiss, M.; Isshiki, T.; , "Cool MPSoC programming," Design, Automation & Test in Europe Conference & Exhibition (DATE), 2010 , vol., no., pp.1488-1493, 8-12 March 2010.

[HSCTB10] Kai Huang; Santinelli, L.; Jian-Jia Chen; Thiele, L.; Buttazzo, G.C.; , "Periodic power management schemes for real-time event streams," Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on , vol., no., pp.6224-6231, 15-18 Dec. 2009.

[LPFMT10] Paul Lokuciejewski, Sascha Plazar, Heiko Falk, Peter Marwedel, Lothar Thiele, "Multi-objective Exploration of Compiler Optimizations for Real-Time Systems," Object-Oriented Real-Time Distributed Computing, IEEE International Symposium on, pp. 115-122, 2010 13th IEEE International Symposium on Object/Component/Service-Oriented Real-Time Distributed Computing, 2010.

Linköping - DTU

[IPEP2010a] V. Izosimov, P. Pop, P. Eles, Z. Peng, Synthesis of Flexible Fault-Tolerant Schedules for Embedded Systems with Soft and Hard Timing Constraints, in Design and Test Technology for Dependable Systems-on-chip, Raimund Ubar, Jaan Raik, Heinrich Theodor Vierhaus Eds., Information Science Publishing, 2010.

[IPEP2010b] V. Izosimov, P. Pop, P. Eles, Z. Peng, Value-Based Scheduling of Distributed Fault-Tolerant Real-Time Systems with Soft and Hard Timing Constraints, IEEE Workshop on Embedded Systems for Real-Time Multimedia, Scottsdale, AZ, USA, October 28-29, 2010.

[IPEP10] V. Izosimov, P. Pop, P. Eles, Z. Peng, "Scheduling and Optimization of Fault-Tolerant Embedded Systems with Transparency/Performance Trade-Offs", in ACM Transactions on Embedded Computing Systems (in press)

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

2.5 Keynotes, Workshops, Tutorials

Keynote : ARTEMIS - Deriving Research Problems from Complex Societal Challenges (Rolf Ernst, TU Braunschweig)
ARTEMIS Summer Camp
Rome, Italy - June 10, 2010

The ARTEMIS Industry Association for R&D actors in embedded systems in Europe is the meeting place where key industry and R&D actors identify topics for major R&D project proposals and form consortia. In his keynote speech at the ARTEMIS Summer Camp 2010, Rolf Ernst discussed how to derive research problems from complex societal challenges and anticipated that the new area of embedded systems-of-systems will be driven by comprehensive societal challenges rather than by individual application domains. In this context robust networks of embedded systems, mixed criticality networks, and autonomous systems requiring new design processes have been identified as resulting scientific challenges.

(see also the report on the Design for Adaptivity activity)

http://www.artemisia-association.org/presentations_2

Keynote : Certification of Trusted MPSoC Platforms (Rolf Ernst, TU Braunschweig)
International Forum on Embedded MPSoC and Multicore, MPSoC'2010
Gifu, Japan - June 29, 2010

The MPSOC event brings together key R&D actors from the different fields required to design embedded Multiprocessor SoC (MPSoC) and Multi-core SoC. In his talk Rolf Ernst presented some of the main challenges regarding system integration on MpSoCs. Safety critical system integration on MpSoCs becomes more difficult due to physical resource sharing and thus may have to face increasing certification costs in comparison to distributed systems. Solutions from an academic point of view were suggested.

http://www.mpsoc-forum.org/2010/agenda.html#Lectures

Keynote : Safety, Efficiency and Autonomy – Mastering Conflicting trends in Embedded Systems Design. (Rolf Ernst, TU Braunschweig)
DIPES Conference 2010
Brisbane, Australia - September 20, 2010

The 7th IFIP Conference on Distributed and Parallel Embedded Systems (DIPES 2010) was hosted as part of the IFIP World Computer Congress WCC 2010. Rolf Ernst gave a keynote presentation on current research results and future trends in the field embedded system design with focus on safety, efficiency and autonomy.

http://wcc2010.com/DIPES2010/index.html

Keynote : Embedded Systems Research for Complex Societal Challenges, (Rolf Ernst, TU Braunschweig)
ESI Symposium 2010

Eindhoven, Netherlands - December 2, 2010

On December 2, 2010, ESI the Embedded Systems Institute located at the university campus in Eindhoven had the third annual symposium. The symposium covered a wide range of applied research topics on embedded systems and was organized together with Point One, an open association of high-tech industry and knowledge institutes with research and development in the Netherlands on nanoelectronics, embedded systems, and mechatronics. In this keynote Rolf Ernst highlighted design challenges due to increasing system complexity and due to requirements for integration and certification of mixed-critical embedded systems. *(another topic covered in this talk is related to the Design for Adaptivity activity)*

http://www.esi.nl/symposium/

Keynote : Smart Distributed Sensors for Adaptive Green Services (Davide Brunelli, UNIBO)
2010 18th IEEE/IFIP International Conference on VLSI and System-on-Chip - VLSI–SoC
Madrid, Spain - 27–29 September 2010

VLSI-SoC 2010 is the 18th in a series of international conferences sponsored by IFIP TC 10 Working Group 10.5, IEEE CEDA and CASS that explores the state-of-the-art and the new developments in the field of Very Large Scale Integration (VLSI), System-on-Chip (SoC) and their designs. Abstract :*Public administrations, enterprises and citizens are increasingly going green and are looking to the information technology as the way to reduce energy consumption and to become more environmentally responsible. In particular reducing buildings overall energy consumption is certainly on of the major effort. Distributed and pervasive sensing and monitoring will play a key role in achieving this goal and will pose novel research challenges in distributed monitoring such as how the space is used, how the data is compressed and how controlling devices are interfaced. The talk will present examples and case studies of architectures for Green Computing leveraging connected sensing systems, networks, and devices.*

http://www.vlsi-soc.com/

Keynote : P2012: Designing many-core platforms for silicon-efficient embedded multimedia computing (Luca Benini, UNIBO)
5th symposium on Computer Architecture and Digital Systems CADS2010
Tehran, Iran - September 23-24, 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://cs.ipm.ac.ir/cads2010/index.jsp

Mini-Keynote : A Bio-Inspired Reconfigurable Hardware Architecture Supporting Self-organisation and Self-healing (Jan Madsen, DTU)
10th International Forum on Embedded MPSoC and Multicore
Gifu, Japan, july 2010

Electronic devices based on modern chip technology are susceptible to both transient and permanent failures due to increased integration and to process variability of the chip technology. In this talk, we will present a reconfigurable hardware platform, which is capable of Self-organisation and Self-healing based on biological principles. The platform is a multicore chip, where each core acts as a cell. An application is compiled into a compact representation

resembling the biological DNA. Self-organisation allows the cells of the platform to autonomously determine their functionality from the DNA based on their position. In case of a cell failure, other cells are able of detect this and re-establish the lost functionality at a nearby idle-cell through the sharing of DNA, effectively obtaining Self-healing of the platform.
http://www.mpsoc-forum.org/2010/index.html

Keynote : Schedule Memory Access, not Threads (Martin Schoeberl, DTU)
10th International Forum on Embedded MPSoC and Multicore
Gifu, Japan - july 2010
Chip-multiprocessing is considered the future path for performance enhancements in computer architecture. Eight processor cores on a single chip are state-of-the art and several hundreds of cores on a single die are expected in the near future. General purpose computing is facing the challenge how to use the many cores. However, in embedded real-time systems thread-level parallelism is naturally used. We assume a system where we can dedicate a single core for each thread. In that case classic real-time scheduling disappears. However, the threads, running on their dedicated core, still compete for a shared resource, the main memory. A time-sliced memory arbiter is used to avoid timing influences between threads. The schedule of the arbiter is integrated into the worst-case execution time (WCET) analysis. The WCET results are used as a feedback to regenerate the arbiter schedule. Therefore, we schedule memory access instead of CPU time.
http://www.mpsoc-forum.org/2010/index.html

Tutorial : Hardware/Software Codesign of Embedded Systems using GEZEL (Jan Madsen and Gilberto F. Marchioro, DTU)
28th Norchip Conference
*Tampere, Finland – 14.11.*2010
This tutorial presents a model-driven approach to hardware/software codesign based on the open-source GEZEL framework (language and tools) from Virginia Tech. The GEZEL language is a cycle-based hardware description language based on the Finite-State-Machine with Datapath (FSMD) model. The GEZEL tools offer stand-alone hardware simulation, cosimulation, and code-generation into synthesizable (VHDL) code. GEZEL allows for the inclusion of new userdefined cosimulation interfaces, through user-defined library-block extensions in C++, which allows the modeling and simulation of IP-blocks.
The tutorial will cover the basics of the GEZEL language and how to build cosimulation systems based on ARM cores. Recent work at DTU Informatics on how to extend the GEZEL framework with support for formal hardware verification will be presented. Finally, the usage of GEZEL for an Embedded Systems course targeting bachelor software students at their third semester, will be addressed.

Seminar: Wireless Sensor Networks (Jan Madsen, DTU)
Special Interest Group on Green-IT, Infinit innovation network on ICT
*Lyngby, Denmark – 17.11.*2010
DTU organized the a seminar on wireless sensor networks, with focus on how this technology can help monitoring climate and environmental changes, and monitoring the “health” of building constructions. The event had 24 participants, where half were from industry. One of the presentations were given by Jan Buetel from ETHZ.
http://www.infinit.dk/dk/nyheder_og_arrangementer/arrangementer/afsluttede_aktiviteter/20101011_01.htm

Tutorial : R. DAVID, CEA LIST, Multi-core architectures design space analysis for embedded systems
RSP 2010 : 21st IEEE International Symposium on Rapid System Prototyping
Firfax, USA

Tutorial: Embedded Systems and their Physical Environment – Sensor Networks in Environmental Applications: A key Application
DATE 2010

Dresden, Germany – 9.3. – 11.3.2010

The tutorial was organized by partner EPFL (Giovanni DeMicheli). The talk presented joint results with UNIBO on energy harvesting and their use in sensor network applications in harsh environments.

Tutorial: Embedded Software for SoC Design
ASPDAC 2010

Taipeh, Taiwan – 18.1. – 21.1.2010

For many emerging embedded applications, high performance is required: High-quality multimedia processing in consumer electronics, software defined radio in communications systems, or real-time diagnostics in medical systems are typical examples. A frequent choice for digital signal processing systems will be heterogeneous multiprocessor system-on-chip (MPSoCs) because of their computational power, programmability, and low power dissipation. Software development plays a central role in handling the increasing complexity of applications implemented on MPSoCs. Productively programming heterogeneous MPSoCs requires support for concurrency, timing, heterogeneity, scalability, and hardware/software system integration. This support is only provided to a very limited degree by the traditional practice of using C/C++ and a board support package to program single and multiprocessor signal processing systems. Recognizing this challenge, a variety of techniques and complete software design flows have been proposed that shift the software development for MPSoCs to higher levels of abstraction. In those design flows, applications are developed using a high-level application programming interface (API), or a model of computation. These approaches attempt to assist software developers in the necessary high-level design decisions and allow automating certain steps in the design flow. A particular challenge is the performance analysis of MPSoCs which is required for early design space exploration and final system verification. Simulation-based methods are not well suited for this purpose due to long run-times and missing corner-case coverage. To overcome these limitations, formal performance analysis methods that scale to large systems and provide guarantees for meeting real-time constraints have been developed. Embedding formal performance analysis into the MPSoC design cycle requires the generation of a faithful analysis model and its calibration with the system-specific parameters.
The tutorial provided an overview about the major challenges in multiprocessor software development. We will present a taxonomy of software design flows based on this analysis, review current MPSoC software design flows and classify them based on the associated challenges. As an example, a design flow is presented that integrates a modular performance analysis method into the MPSoC programming environment. The result is an MPSoC software design flow that allows to automatically generating the system implementation together with an analysis model for system verification.

Summer School: Models for Embedded Signal Processing Systems.
Leiden University and Lorentz Center
Leiden, The Netherlands – 30.8.2010 - 3.9.2010

During the system level design process of an embedded system, a designer is typically faced with questions such as whether the timing properties of a certain system design will meet the design requirements, what architectural element will act as a bottleneck, or what the on-chip memory requirements will be. Consequently it becomes one of the major challenges in the design process to analyze specific characteristics of a system design, such as end-to-end delays, buffer requirements, or throughput in an early design stage, to support making important design decisions before much time is invested in detailed implementations. This

analysis is generally referred to as system level performance analysis. If the results of such an analysis are able to give guarantees on the overall system behavior, it can also be applied after the implementation phase in order to verify critical system properties.
One of the major requirements for models, methods and tools is their support for a modular, component-based design. This aspect covers as well the composition of the underlying hardware platform as well as the software design. Because of the import role of resource interaction, these components not only need to talk about functional properties but also about resource interaction.
The course covered the following aspects of system level performance analysis of distributed embedded systems: Approaches to system-level performance analysis. Requirements in terms of accuracy, scalability, composability and modularity. Modular Performance Analysis (MPA): basic principles, methods and tool support. Examples that show the applicability: An environment to map applications onto multiprocessor platforms including specification, simulation, performance evaluation and mapping of distributed algorithms; analysis of memory access and I\O interaction on shared busses in multi-core systems.

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 allocated to outline safety requirements and to present the impact of fault tolerance and fail-safe mechanisms on real-time systems properties. Possible solutions for the analysis and optimization of mixed-critical systems were discussed. *(Further topics addressed in this lecture are related to the activity report on Platform and MPSoC Analysis and on the report on Design for Adaptivity Activity)*

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

Presentation: Challenges of Mapping Real-Time Streaming Applications to General-Purpose Many-cores
ArtistDesign Workshop on Mapping Applications to MPSoCs 2010
Schloss Rheinfels, St. Goar, Germany, June 29-30, 2010

Jonas Diemer (TU Braunschweig) gave a presentation on the mapping of real-time applications in a general-purpose system with minimal support for quality-of-service.

http://www.artist-embedded.org/artist/Program,1822.html

Presentation: Predictablity in General-Purpose Many-Cores
Intel® European Research & Innovation Conference, ERIC 2010
Braunschweig, Germany, September 21-22, 2010

At the ERIC Conference, a Poster presentation was given by Jonas Diemer (TU Braunschweig) on novel network-on-chip architectures that enable future general-purpose many-core to efficiently execute real-time applications with guaranteed timing requirements. This topic was also presented at the parallel 10-year Jubilee of Intel Braunschweig, who funded the corresponding project COMPOSE.

http://www.intel.com/corporate/education/emea/event/irc/deu/

Invited Talk: Mastering Mixed Criticality - A Growing Challenge in Embedded Systems Integration (Rolf Ernst, TU Braunschweig)
Symtavision News Conference 2010
Braunschweig, Germany – September 29 - 30, 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. Rolf Ernst was invited to present current research results to this audience.

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

Invited Paper : R. DAVID, CEA LIST, Resources management in multi- and many-cores architectures: Issues and trends, invited paper, VLSI-SOC Conference, Madrid, Spain, 2010

3. Milestones, and Future Evolution

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

In 2011**, CEA LIST** and **UNIBO** will continue to dynamically manage computing and memory resources. The basic principle implemented in 2010 will be extended to optimized furthermore the reactivness of the system. Also, **CEA LIST** and **UNIBO** will try to extend this hardware support so that to be able to dynamically manage resources in massively parallel device, preserving a shred memory space but non uniform memory accesses.

EPFL and **UNIBO** will continue in year 4 with particular reference to NOCs and 3D Integration.

UNIBO will continue its focus on GPGPUs starting the deveopment of a more complex environment which permit to enenache the feature of the GPU. Additionally **UNIBO** will continue on the research on Energy Harvesting exploiting the hardware and software capabilities of a multi-source enrgy harvest.

UNIBO and **SSSA** will continue on the research activity about scheduling problem on multi-core systems. We expected futrther improvements on refining Elastic Scheduling algorithms. UNIBO will continue to simulate on a TDMA Bus and will addrees other BUS configurations. The goal is to ensure the highest utilization of the multi-core system and to develop efficient scheduling of multi-task applications with inter-task dependencies.

TU Braunschweig will continue the work on the reliability analysis methods. Approaches currently under development will be compared with respect to computational complexity, accuracy and applicability for practical problems. In the field of multi-core analysis, the pool of available modeling techniques will be extended to incorporate different hardware and software recovery mechanisms with special emphasis on mixed-criticality, including potential design trade-offs to enable cost- and power-efficient designs combined with safety guarantees.

In year 4 of ArtistDesign, **ETHZ** will continue to jointly work on new methods to map algorithms onto scalable MPSoC platforms. We will concentrate on more adaptive strategies that allow for dynamic applications as well as changing requirements. At the same time, we will take temperature constraints into account in the analysis.

Linköping will continue the research concerning temperature aware and energy efficient design of real-time embedded systems. During the fourth year we will continue to elaborate new techniques for temperature aware power management. One of the goals is to extend our temperature aware idle time distribution technique such that it can be applied on-line for the management of dynamic slack.

DTU will, in 2011, extend Tinuso to include interfaces to on-chip networks and support for multiple programming models. The core will also be extended to support floating point arithmetic.

The collaboration with IBM on the aforementioned parallel programming tools will be further extended in 2011. The usefulness of the tool will be evaluated using realistic workloads studied within the SMECY project.

The work on data structures for operating systems will be extended and more fundamental data structures will be explored. As part of the SMECY project a new distributed light weight operating system for many-core architectures will be developed in 2011. The fundamental design principles of the system is that the per core resources are scarce and that the organization of the operating system should follow the partitioning of the many-core architecture.

DTU and **Linkoping** will continue the research on fault tolerance in distributed systems. The Research Agenda for Mixed-Criticality Systems 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 the avionics area, 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. **DTU** plans to tackle the problem optimising the time-partitions for mixed-criticality real-time distributed embedded systems. Linkoping will provide extensions of their schedulability analysis for mixed time/event-triggered real-time applications, to consider the impact of time-partitions on the worst-case response time of event-triggered tasks scheduled using fixed-priority preemptive scheduling.

During the synthesis of DMBs, researchers have assumed that the biochemical operations are executed on rectangular virtual devices, formed by grouping several adjacent electrodes. However, the operations can actually execute by routing the droplets on any sequence of electrodes on the array. Hence, DTU plans to propose a model for biochemical operations that will eliminate the concept of virtual modules and allow the droplets to move on the chip on any route during operation execution. The synthesis problem will thus be transformed into a routing problem.

DTU will continue its research on biochips. Regarding Continous Fluidic Biochips (CFBs), DTU plans to tackle the optimization of how a biochemical application is performed on a biochip. DTU will look into cell culture biochips, where several cell colonies are exposed to soluble compounds and monitored in real-time to determine the right combination of factors that lead to the desired results. The application considered are a full-factorial experimental design, where all possible combinations of compounds are applied. DTU is interested to synthesize the settings of an experimental design such that the throughput is maximized.

DTU will continue its research on dynamically reconfigurable platforms. The focus will be on efficient techniques to detect faults and initiate the self-healing process. The full FTS application from NASAs JPL will be tested on the eDNA platform.

-- Changes wrt Y2 deliverable --

For several partners, the main objective area to tackle in Y4 are the same of Y3. Some additional research topic has been added to the current plans on Y4.

3.2 Current and Future Milestones

- **Modelling of and evaluation of fault-tolerance mechanisms with respect to real-time**

  Y3: In 2010 these prototypes will be refined. Further on the consideration of other protocols will be taken into account, especially FlexRay as an up-to-date automotive bus system.

  This milestone has been partially achieved. The research on reliability analysis has been intensified in different directions to incorporate multi-cores analysis as well as the effects of clock drift on reliability estimation for example. FlexRay has not been considered in detail, instead the focus has been set on system analysis.

Y4: In the fourth year, the concept of system reliability analysis will be refined. The goal is to compare different analysis approaches with respect to computational complexity, accuracy and applicability for practical problems. Modeling techniques in the field of multi-core analysis will be extended with special emphasis on mixed-criticality.

- **Predictability and modularity of MPSoC for Real-Time applications.**

  Y2: Erika RTOS porting on MPARM (UNIBO) The kernel provides all functionalities to specify and run such tasks.

  Y3: Designing a system where Elastic Scheduling and the TDMA Bus work synergistically to ensure the highest utilization.

  These milestones has been fulfilled

  Y4: In the next year, the goal is to ensure the highest utilization of the multi-core system and to develop efficient scheduling of multi-task applications with inter-task dependencies by exploring several configurations and changing parameters (e.g. low power aging

- **MPSoC architectures and programming models**

  Y4: New models and methods for mapping of algorithms onto MPSoC platforms defined, focussing on dynamically changing applications. Taking into account upper bounds on system temperature in MPSoC software design.

- **Design of fault tolerant distributed embedded systems**

  During the third year Linköping and DTU will continue their work on fault tolerant distributed and multiprocessor embedded systems. The emphasis during the third year will be on the optimised implementation of error detection techniques.

  The above milestone has been fulfilled.

  Y4: Fault aware hardware software partitioning for systems on chip.

- **Energy efficient embedded system design**

  Linköping will continue the research concerning temperature aware and energy efficient design of real-time embedded systems. During the third year we will continue our work on elaboration of fast and sufficiently accurate analytical temperature models for the system level. We will also elaborate, in addition to our temperature aware dynamic voltage scaling techniques, new techniques for temperature aware power management.

  The above milestone has been fulfilled.

  Y4: Extend our temperature aware idle time distribution technique such that it can be applied on-line for the management of dynamic slack.

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:

CEA LIST

- **MC2H (Many-core for Computing and Healing). French R&D cooperation program (Nano 2012**
- **SCALOPES (SCAlable LOw Power Embedded platformS). ARTEMIS project**

EPFL

PRO3D
De Micheli's activities will be funded in part by the FP7 PROD3D grant – together with partners Thiele, Benini, Sifakis and Jerraya (amog others). They will also be funded through an ERC senior grant.

TU Braunschweig

- **COMPOSE Project**
  Future computer architectures are likely to have tens to hundreds of processor cores, running different application classes in parallel. This project investigates techniques to combine their heterogeneous requirements and make such architectures predictable. Examples are partitioning of on-chip memories, quality of service for the network-onchip, and mechanisms for flexible yet efficient data transfers.
  http://www.ida.ing.tu-bs.de/en/research/projects/compose
- **RECOMP Project**
  The increasing demand for processing power poses new challenges on the design of modern embedded systems. The adoption of multi-core processors seems to be a promising approach to tackle these challenges and to achieve further performance improvements combined with a reduction of energy consumption. However using these architectures in safety-critical applications such as aerospace, automotive, health or industrial automation leads to additional requirements, because such systems have to be certified according to domain specific safety-standards. Driven by these new requirements the ARTEMIS project RECOMP (Reduced Certification Costs for Trusted Multi-core Platforms) aims at establishing methods, tools and platforms for cost-efficient certification and re-certification of safety-critical multi-core systems. Special emphasis is placed on the consideration of mixed-criticality, i.e. systems containing both safety-critical and non-safety-critical components.
  http://www.ida.ing.tu-bs.de/en/research/projects/recomp
- **ASTEROID**
  The goal of ASTEROID is the development of new hardware and software concepts to enhance reliability of safety-critical embedded systems. Special emphasis is put on the appliance of microkernels, in particular L4. The project is funded by the German research funding organization DFG (Deutsche Forschungsgemeinschaft).
- **Toyota-ITC**
  In 2008 a research cooperation between TU Braunschweig, Toyota Information Technology Center (T-ITC), and Symtavision GmbH has been initiated, investigating the effects of errors on reliability and safety of real-time networks. Follow-up project have been started in 2009 and 2010.

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)
- MICS Mobile Information and Communication Systems (Swiss National Science Foundation)

UNIBO

- ICT-Project PREDATOR (EU FP7)
- ICT-Project GALAXY (EU FP7)
- PROD3D Programming for Future 3D Architecture with Many Cores (EU FP7)
- ICT-Project Scalopes (Artemis JU)
- Industrial funding on Sensor Networks from Telecom Italia spa

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

- **DaNES** (Danish Network for Embedded Systems, funded by the Danish Advanced Technology Foundation), Denmark. Period 2007-2010.
- **MoDES** (project on Model Driven Development of Intelligent Embedded Systems funded by the Danish Strategic Research Council), Denmark. Period 2006-2009.
- **ProCell** (project on programmable cell chip: culturing and manipulation of living cells with real-time reaction monitoring funded by the Danish Strategic Research Council),
- **SYSMODEL** (System-Level Modeling for SMEs) funded by ARTEMIS JU. Period 2009-2011.
- **RECOMP** funded by ARTEMIS JU. Period 20010-2013.
- **SMECY** funded by ARTEMIS JU. Period 20010-2013.
- **ASAM** funded by ARTEMIS JU. Period 20010-2013.

-- Changes wrt Y2 deliverable --

Some EU FP7 projects has been added to contribute to the NoE (e.g. PROD3D Programming for Future 3D Architecture with Many Cores)

4. Internal Reviewers for this Deliverable

- **Dr. Maja D'Hondt** (Imec vzw)
- **Prof. Giorgio Buttazzo** (SSSA, Pisa, Italy)