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

Activity Progress Report for Year 4

# Platform and MPSoC Design

Cluster:

**Hardware Platforms and MPSoC**

Activity Leader:

**Prof. Luca Benini (UNIBO)**

**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 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 | January 20th 2012 |

# 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: Dr. Chantal Ykman-Couvreur – 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 Y3 deliverable --*

*Dr. Chantal Ykman-Couvreur replaces Dr. Maya Dhondt  as IMEC representative*

---

## 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. 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. Christian Gries, Intel (Germany)
>  Thermal control and energy management for multi-core platforms

Dr Marek Jersak Symtavision (Germainy)
>  Design for predictable and composable MpSoCs and runtime prediction tools

---

*-- Changes wrt Y3 deliverable --*

---

*Dr. Valter Bella from Telecom Italy has not participated to the activities and has been removed.*

---

## 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 Y3 deliverable --*

*No changes with respect to Year 3.*

## 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 Y3 deliverable --*

*No changes with respect to Year 3.*

## 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  collaborated on dynamic memory management optimizations at the system level for single processor 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 are collaborating 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 Y3 deliverable --*

*No changes with respect to Year 3.*

---

## 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 Y3 deliverable --*

*No changes with respect to Year 3.*

---

## 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 Y3 deliverable --*

*This section was already presented in the Y3 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 Y3 deliverable --*

*This section was already presented in the Y3 deliverable, in sections 1.8.*

---

## 1.9   *Work achieved in Year 3*          *(Jan-Dec 2010)*

In Y3 **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. On the one hand 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 after 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, t**he **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.

---

*-- No changes wrt Y3 deliverable --*

*This section was already presented in the Y3 deliverable, in sections 1.9 and 3.1.*

---

## 1.10 Problems Tackled in Year 4          *(Jan-Dec 2011)*

**TU Braunschweig** has continued the work on *reliability analysis for unreliable components in safety-critical systems* with its industrial partners (in particular Symtavision and Toyota-ITC). Additionally the FP7 project CERTAINTY has been started, focusing on composable and scalable analysis and design methods for mixed-criticality systems.

Concerning the cooperation with Symtavision and Toyota-ITC special emphasis has been put on the challenges of integrating reliability analysis into industrial design and verification tools such as Symtavisions SymTA/S toolsuite. In this context especially the scalability of reliability analysis has been pointed out as a crucial issue. It was a major objective of the research in 2011 to investigate the applicability of reliability analysis to real-world CAN bus examples and to optimize the analysis algorithms in a suitable way such the industrial-driven requirement on performance and scalability are met. For this purpose different approaches have been considered, from simple implementation issues up to highly efficient approximation methods which generate conservative reliability estimations with adjustable granularity and accuracy.

In the ASTEROID project which started in October 2010, TU Braunschweig together with TU Dresden focused on designing, modeling and analysis of novel fault-tolerance mechanisms for mixed-critical workloads. In such application scenarios, tasks of different criticality (e.g. safety critical and best effort) are executed on the same MPSoC. The major challenge of such systems is to employ an optimal mixture of fault-tolerance mechanisms to provide a low degree of over-provisioning while maintaining the safety-requirements of the critical tasks.

In the CERTAINTY project which started in November 2011, the main objective of TU Braunschweig is to develop new methods for reliability analysis on system level by taking different integration issues into account. These issues include physical composition effects caused by correlated error models as well as logical composition effects due to resource sharing by applications with different time- and safety-criticality. TU Braunschweig is currently involved in refining the actual problem statement and works on identifying the concrete research problems to be tacked in 2012. This happens in cooperation with industrial project partners such as Thales, which contribute their domain specific requirements on hardware and software.


**ETHZ** has been concentrating on the following three issues:

- Mapping of algorithms onto MPSoC structures. In a joint paper, several approaches have been compared, see [MTKBTHLXH11].
- Considering thmeral issues in the desihgn of MPSoC systems, see [TSYB11].
- Developing methods that allow an efficient resource isolation in MPSoCs, see [KCT11].

As can be seen from the above results, there has been substantial progress in combining various methods and approaches. In addition to last year, we have been able to also handle resource isolation and adaptive control in the context of designing predictable embedded multiprocessor systems.


In 2011**, CEA LIST** continued the work started with **UNIBO** on the dynamic management of computing and memory resources. The low cost hardware synchronizer component designed in 2010 was refined and transferred to partner STMicroelectronics for being integrated into a chip designed in the *Platform 2012* project. This hardware synchronizer component was also used as a basis for the specifications and the design of a hierarchical control infrastructure for the *COBRA CATRENE* project. In this context, **CEA LIST** investigated a two-level control infrastructure based on a global and centralized programmable controller and local cluster-level controllers taking advantage of the acceleration provided by the hardware synchronizing component. Based on this hierarchical infrastructure, both a hierarchical vision of the software and the appropriate hardware acceleration for supporting migration and preemption at the global level are currently studied. Additionnaly, **CEA LIST** also participated to the Locomotiv chip, a 32nm implementation of a multi-chip *Platform 2012* instance. This multi-chip architecture is controlled by a global controller sub-system implemented on an FPGA. One aim

of this platform is to illustrate the different mechanisms used for dynamically managing resources in massively parallel device

**CEA LIST** has also continued its work on its SCMP asymmetric multicore plartform by transferring the FPGA implementation of SCMP on an emulation platform. Based on this work, investigations were started in mid-2011 in order to link this emulation platform to the **CEA LIST** SESAM virtual prototyping platform in order to allow the possibility to refine a prototyped design by integrating emulated hardware blocks in the architecture rather than their virtual prototypes.


**Linköping** has continued the research concerning temperature aware and energy efficient design of real-time embedded systems. The emphasis was on In the issue of leakage energy optimization through temperature aware idle time distribution (ITD). We have proposed an on-line ITD technique for leakage energy consumption minimization, where both static and dynamic idle time are considered.

Furthermore the **Linköping** group has considered distributed embedded systems in which privacy or confidentiality of the internal communication is critical, and has proposed an approach to optimizing cryptographic algorithms under strict timing constraints. A technique has been developed to search for the best system-affordable cryptographic protection for the messages transmitted over the internal communication bus of a distributed embedded system.


**EPFL** together with **UNIBO** have worked on further development of the synthesis tool for designing application specific Network-on-Chip (NoC) for 2-D and 3-D systems especially with tight constraints on the power consumption [SMBM11]. NoCs can provide a salable interconnect platform for 3-D integrated chips but require tools for topology design and exploration in order to compare 2-D and 3-D NoC-based systems [SMBM11]. In this work, further challenges are presented by 3-D integration as it is expensive to have a global synchronous clock in 3-D chips and the global interconnect also has to cope with multiple frequency and voltage islands.

In addition, many Systems-on-Chip (SoC) require a certain levels of Quality-of-Service. For some services (e.g. interrupts) providing worst case guarantees is crucial for the SoC to function correctly. As presented in [SRMBM11], with an adequate design of the system NoC topology worst case guarantees can be provided without any specialized hardware requirement.

On the other hand, current SoCs are DRAM centric. As a consequence the external DRAM memory can become a bottleneck in complex NoC-based system, and packets queuing up in the network can affect the performance of non DRAM traffic. One solution is having separate NoC channels for DRAM and non DRAM traffic. Such a solution also presents other benefits in case where the core data size is heterogeneous [SMBM11b].

Furthermore, concerning the design of 3-D for high performance ICs, EPLF has worked as well on techniques to design 3-D monolithic ICs, developing CELONCEL, a placement tool targeting the generation of high quality 3-D layout [BCTBEFPM11][BCTBPM11].


**DTU** continued its work on architectures and programming models for MPSoC. During 2011, DTU has participated in several ARTEMIS projects, SMECY, RECOMB and ASAM. One focus, related to ASAM has been to address the initial configuration of an MPSoC consisting of configurable ASIP cores. Given an application that has been parallelized, what will be the most likely optimized platform (number of cores and interconnection topology), which can reach the requirements after detailed ASIP optimization at a later stage. This work is done in collaboration with ST, Intel, TU Eindhoven and TUBS.

In 2011, **DTU** has continued its participation in the SMECY project. The work has focused on architectures, operating systems and programming models for MPSoCs. DTU has also further strengthen collaborations with IBM Haifa.

**DTU** has in 2011 further developed the eDNA prototype of a new reconfigurable hardware platform with self-organizing and self-healing capabilities. The eDNA technology was successfully tested in a space application together with NASA JPL in 2010 and is now being patented in Europe, USA and Japan. Finally, DTU is exploring possibilities of making a spin-off company.

**Linkoping** and **DTU** have continued their work on fault tolerant distributed and multiprocessor embedded systems. **DTU** has focused on mixed-criticality hard real-time applications on a given heterogeneous distributed architecture. Applications have different criticality levels, captured by their Safety-Integrity Level (SIL), and are scheduled using static-cyclic scheduling. We have also considered fixed-priority preemptive tasks (event-triggered), where we have extended the response-time analysis from **Linkoping**, which considers time-triggered tasks that can interfere with the event-triggered tasks.

**IMEC** has been concentrating on the following issues: i) **Run-time monitoring of application parameters**: for many applications, the processing requirements to obtain results at real time can be impractical. This is due to the increasing complexity of the applications developed for advanced platforms. Nevertheless, a lower output quality might be acceptable if the results are obtained within the time required. In order to trade the output quality for the requirements (e.g., application deadline, power budget of the platform) to obtain the results, a set of well-chosen application parameters must be tuned/monitored at run time once the requirements are set by the user. This technique is developed within the European project 2PARMA, about parallel paradigms and run-time management for many-core architectures. ii) **Distributed and hierachical run-time resource management** for audio-driven surveillance applications on heterogeneous MP-SoC platform. This work is performed within the European project COMPLEX, aboutr codesign and power mangement in platform-based design space exploration. iii) **Develop monitoring and control techniques** in the middleware of the system to automatically adapt platform services to application requirements. This work is performed within the European project Pharaon, about parallel and heterogeneous architecture for real-time applications

*-- The above is new material, not present in the Y3 deliverable --*

# 2. Summary of Activity Progress in Year 4 *(Jan-Dec 2011)*

## *2.1 Technical Achievements*

**Reliability analysis with mixed-criticality workloads (TU Braunschweig, Toyota-ITC, Symtavision, CERTAINTY Project Consortium, TU Dresden)**

The process of industrial integration has been advanced successfully. Feasibility of reliability analysis for real-world systems could be demonstrated based on different CAN bus examples provided by Toyota-ITC. In this context analysis performance could be improved by orders of magnitude compared to former implementations used for internal research purposes only.

Especially the adjustable analysis granularity has been pointed out as an effective measure to accelerate reliability analysis. In this context nearly no loss of accuracy could be detected. Granularity is adapted automatically during analysis, such that reliability analysis is basically executed with an optimal accuracy-performance tradeoff. Most of the algorithmic improvements have been implemented within a prototype plugin for SymTA/S, which is going to be evaluated by our industrial partners by now.

A joint publication together with Symtavision summarizing the results on reliability and safety analysis for CAN buses and multi-core over the last years has been presented at the SAE World Congress 2011 [SAEFJ11]. Results of former research activities on error models and the effect of error correlations on reliability have been published on PRDC 2011 [SAE2011].

In collaboration with TU Dresden we designed novel fault-tolerance hardware and software features targeting at the Fiasco L4 Microkernel. This comprises application level redundancy and task restart/rollback on software level as well as a fingerprint based hardware voting mechanism which was also implemented on a FPGA demonstrator.

The timing effects of redundancy and re-execution of embedded real-time software were modeled and analyzed using techniques known from compositional performance analysis. Results have been published at ESWEEK 2011 [ASE11].

**RECOMP-Project: Integrated Dependable Architecture for Many Cores (IDAMC)**

TU Braunschweig develops a many-core architecture for mixed-safety-critical applications within the RECOMP Project. This architecture shall enable the integration of applications with different levels of criticality. It is centered around a special network-on-chip [DE10a], which offers data transfers with predictable timing. This is based on the work in the COMPOSE Project which was finished in 2010. This NoC interconnects different "tiles" which contain general-purpose processors based on the LEON3 SPARC architecture. For this purpose, a network-interface is developed within RECOMP which facilitates advanced functionality for virtualization and monitoring. RECOMP also includes the analysis of the timing and reliability properties of the platform [DRNSE11, SAEFJ11, ASE11].

**Comparting mapping strategies to MPSoC architectures (ETH Zurich, TU Dortmund, U Erlangen Nuremberg)**

The advent of embedded many-core architectures results in the need to come up with techniques for mapping embedded applications onto such architectures. The results described in [MTKBTHLXH11] are a representative set of such techniques. The techniques focus on optimizing performance, temperature distribution, reliability and fault tolerance for various

models. Generated mappings have to include solutions to the scheduling, resource allocation and resource assignment problems. The type and number of resources (processors, buses, etc.) may be either be already specified in the initial input by the designer or they may be generated during the design steps. If this information is not provided by the designer, it has to be generated by solving the resource allocation problem. Resource assignment is supposed to map computations and communication to hardware resources. Solutions to the scheduling problem provide a mapping from operations to the times during which these operations are performed. Multiprocessor scheduling will sometimes already indicate the hardware resources to be used, i.e. the resource assignment problem is included in the scheduling problem. Resource assignment algorithms can be of various types. For example, computations may be permanently mapped to hardware resources.

### Thermal Aware Design of MPSoCs (ETH Zurich)

Nowadays, the reliability and performance of modern embedded multi-processor systems is threaten by the everincreasing power densities in integrated circuits, and a new additional goal of software synthesis is to reduce the peak temperature of the system. However, in order to perform thermal-aware mapping optimization, the timing and thermal characteristics of every candidate mapping have to be analyzed. While the task of analyzing timing characteristics of design alternatives has been extensively investigated in recent years, there is still a lack of methods for accurate and fast thermal analysis. In order to obtain desired evaluation times, the system has to be simulated at a high abstraction level. This often results in a loss of accuracy, mainly due to missing knowledge of system's characteristics. The new results described in [TSYB11] address this challenge and presents methods to automatically calibrate high-level thermal evaluation methods. Furthermore, the viability of the methods for automated model calibration is illustrated by means of a novel high-level thermal evaluation method.

### Resource Isolation in MPSoCs (ETH Zurich)

Servers have been proposed to implement resource reserva- tions on shared resources in Multiprocessor Systems on a Chip. Such reservations isolate the temporal behavior of tasks sharing the shared resources, thereby providing performance guarantees to tasks independent of other tasks. In existing work, resource reservation has been synonymous to utilization (also called bandwidth) on the resource, i.e., we can reserve only a constant fraction of the resource utilization via a server. Such reservation schemes are not suited to serve interrupt-like tasks: tasks that occur seldom but require quick service or tasks with jitter. With this motivation, we presented in [KCT11]  a generalized server algorithm, called Demand Bound Server (DBS), whose offered service is characterized by the demand bound function (dbf) of the task it serves. We show that schedulability of DBS tightly follows that of EDF, and if schedulable a DBS provides a performance guarantee as requested by the dbf of the task..

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

Following the work related to the **Hardware Synchronizer** (HWS) that was started with **UNIBO**, **CEA LIST** adapted this runtime hardware support by including features for asymmetric architectures. This adaptation was necessary for particular *Platform 2012* cluster templates including a programmable controller sub-system called Cluster Controller and it was of great interest for partner **STMicroelectronics**. A dedicated link was introduced between the Cluster Controller (CC) and the Hardware Synchronizer so that the CC can be wakened up by the previously described programmable notifier. This link helps implementing very hierarchical programming models where the Cluster Controller prepares tasks to be executed on the

processing elements of the Cluster. This version of the HWS was then transferred to partner **STMicroelectronics** for being integrated into a multi-cluster chip designed in the *Platform 2012* project. **CEA LIST** also implemented the so-called Asynchronous Reactive Task Management execution engine which is a great improvement compared to the simple reactive task management that has been implemented in cooperation with **UNIBO** in 2010. It can improve the use of computing resources since the Reactive Task Management master resource which launches computing jobs on slaves can now compute part of the job itself. This aRTM execution engine was used in the parallel implementation of a pedestrian detection application together with the cooperative multithreading one.

**CEA LIST** also contributed to the hierarchical control infrastructure of the computing fabric studied in the *COBRA CATRENE* project which involves both **CEA LIST** and **STMicroelectronics**. **CEA LIST** specified a 2-level control infrastructure composed of a programmable global centralized controller and local cluster-level controlling resources. The centralized controller (i.e. Fabric Controller or FC) is based on the STxP70 customizable processor from partner **STMicroelectronics** together with both its associated hardware mailbox based on the **CEA LIST** HWS and its control registers. At cluster level, the local control infrastructure is based on a Cluster Controller sub-system also based on a STxP70 processor associated with the HWS which accelerates both the synchronization and messaging between the computing resources but also the control communications between the Cluster Controller and the computing part of the cluster. After the specifications phase, the Fabric Controller was designed at **CEA LIST** and then transferred to partner **STMicroelectronics**.

Having in mind the hierarchical nature of the *COBRA* Fabric, **CEA LIST** started in mid-2011 to study a hierarchical vision of the software and its associated hardware support. In order to support both performance demanding applications and runtime flexibility, **CEA LIST** proposed to consider an application as a group of tasks at the FC level, so that groups of tasks can be placed and executed on groups of computing resources (i.e. clusters). In order to support flexibility and dynamicity while keeping runtime performance, an hardware support for group preemption and migration is currently under study. Its goal is to help the FC in order to decrease the cost of such operations, which may be required to meet applications deadlines.

**CEA LIST** also participated to an 32nm implementation of a multi-chip *Platform 2012* instance. This multi-chip architecture, called Locomotiv, is controlled by the previously described FC sub-system. CEA LIST implemented this control sub-system on an FPGA motherboard. One aim of this multi-chip platform is to illustrate the different mechanisms used for dynamically managing resources in massively parallel device. Thus, the **CEA LIST** runtime software is being ported on this architecture.

2011 was also the occasion to continue the work on **CEA LIST** SCMP asymmetric multicore plartform. Already implemented on FPGA during 2010, a complete SCMP architecture was implemented on an emulation platform during 2011. The prototype is composed of four LEON3 processors as computing resources with a dedicated controller implemented as a **CEA LIST** AntX processor together with dedicated accelerators for improving control performance and reactiveness. Including its programming model, the support of dataflow execution model and a small OS kernel, the implementation of the SCMP architecture showed interesting results in both application execution acceleration (x2.81) and energy consumption reduction (~33% less) thanks to dynamic power management strategies. Based on this work, investigations were started in mid-2011 in order to link this emulation platform based on EVE solutions to the **CEA LIST** SESAM virtual prototyping platform in order to allow the possibility to refine a prototyped design by integrating emulated hardware blocks in the architecture rather than their virtual prototypes.

## Temperature Aware System-level Power Optimization (Linköping)

**Linköping** has continued the research concerning temperature aware and energy efficient design of real-time embedded systems. At system level, dynamic voltage selection (DVS) is one of the preferred approaches for reducing the overall energy. This technique exploits the available slack times to achieve energy efficiency by reducing the supply voltage and frequency such that the execution of tasks is stretched within their deadline. There are two types of slacks: (1) static slack, which is due to the fact that, when executing at the highest (nominal) voltage level, tasks finish before their deadlines even when executing their worst numbers of cycles (WNC); (2) dynamic slack, due to the fact that most of the time tasks execute less cycles than their WNC. However, very often, not all available slack should or can be exploited and certain amount of slack may still exist after DVS due to the existence of critical voltage. To achieve the optimal energy efficiency, DVS would not execute a task at a voltage lower than the critical one, since, otherwise, the additional static energy consumed due to the longer execution time is larger than the energy saving due to the lowered voltage. 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, the issue of optimizing leakage energy consumption through distribution of idle time is addressed. We proposed an on-line idle time distribution technique for leakage energy optimization which considers both static and dynamic slack. This approach is lookup table (LUT) based and is composed of: an off-line and an on-line phase: (i) The off-line phase prepares a LUT for each task. (ii) At runtime, when a task is finished, the idle time length following the finishing of this task is decided by checking the task's LUT.

## Security aware optimisation of Distributed Embedded Systems (Linköping)

The **Linköping** group has considered distributed embedded systems in which privacy or confidentiality of the internal communication is critical, and has proposed an approach to optimizing cryptographic algorithms under strict timing constraints. A software technique has been developed to search for the best system-affordable cryptographic protection for the messages transmitted over the internal communication bus of a distributed embedded system. Furthermore, we have also proposed an approach to efficiently implement cryptographic algorithms by using hardware/software co-design techniques. The objective is to find the minimal hardware overhead and corresponding process mapping for encryption and decryption tasks of the system, so that the confidentiality requirements for the messages transmitted over the internal communication bus are fulfilled, and time constraints are satisfied.

## 3-D Monolithic Integration and 3D NOCs (EPFL UNIBO)

3-D monolithic integration (3DMI), also termed as sequential integration, is a potential technology for future gigascale circuits. To demonstrate the effect of the 3DMI technology, we compared interconnect-dominated circuit (e.g. LDPC decoder) and gate-dominated circuits [BCTBPM11] showing the tradeoffs in terms of area and wire-length, and overall delay.

Furthermore, we propose two different strategies of stacking standard cells in 3-D without breaking the regularity of the conventional design flow. A placement tool (CELONCEL-placer) targeting the Cell-on-Cell placement problem is proposed to allow high quality 3-D layout generation [BCTBEFPM11]. Our experiments demonstrate the effectiveness of CELONCEL technique, fetching us an area gain of 37.5%, 15.51% reduction in wirelength, and 13.49% improvement in overall delay, compared with 2-D.

In the area of 3D integration, UNIBO and EPFL worked on synthesis tools for 3D NoC systems. Most of the achievements are described in Section 1.10 and in 2.3. The activities performed this year are related to these problems and resulting publications are described in Section 2.4.

**Mixed-criticality systems (DTU, Linköping)**

**DTU** has focused on the optimization of mixed-criticality embedded systems. We were interested to implement mixed-criticality hard real-time applications on a given heterogeneous distributed architecture. Applications have different criticality levels, captured by their Safety-Integrity Level (SIL), and are scheduled using static-cyclic scheduling. We have also considered fixed-priority preemptive tasks (event-triggered), where we have extended the response-time analysis from Linkoping, which considers time-triggered tasks that can interfere with the event-triggered tasks.

Mixed criticality tasks can be integrated onto the same architecture only if there is enough spatial and temporal separation among them. We consider that the separation is provided by partitioning, such that applications run in separate partitions, and each partition is allocated several **time** slots on a processor. Tasks of different SILs can share a partition only if they are all elevated to the highest SIL among them. Such elevation leads to increased development costs. During this year, DTU has focused on modeling the development costs related to certification, and on formulating and solving the following optimization problem.

Given the input mentioned, we were interested to determine (i) the mapping of tasks to processors, (ii) the assignment of tasks to partitions, (iii) the sequence and size of the time slots on each processor and (iv) the schedule tables, such that all the applications are schedulable and the development costs are minimized. We have proposed a Tabu Search-based approach to solve this optimization problem.

**Biochips (DTU)**

**DTU** has looked at continuous flow biochips. Together with Duke University, USA, and the National Cheng Kung University, Taiwan, DTU has investigated the issue of mixed continuous/droplet-based flows for microfluidic biochips. Microfluidic biochips are replacing the conventional biochemical analyzers and are able to integrate the necessary functions for biochemical analysis on-chip. There are several types of microfluidic biochips, each having its advantages and limitations. In this paper we are interested in flow-based biochips, in which the flow of liquid is manipulated using integrated microvalves. By combining several microvalves, more complex units, such as micropumps, switches, mixers, and multiplexers, can be built. Although researchers have pro- posed significant work on the system-level synthesis of droplet-based biochips, which manipulate droplets on a two-dimensional array of electrodes, no research on system-level synthesis of flow-based biochips has been reported so far. The focus has been on application modeling and component-level simulation.

DTU has proposed a system-level modeling and synthesis approach for flow-based biochips. We have developed a topology graph-based model of the biochip architecture, and we have used a sequencing graph to model the biochemical applications. We have considered that the architecture of the biochip is given, and we were interested to synthesize an implementation, consisting of the binding of operations in the application to the functional units of the architecture, the scheduling of operations and the routing and scheduling of the fluid flows, such that the application completion time is minimized. We have proposed a List Scheduling-based heuristic for solving this problem.

**Run-time resource management (DTU)**

**DTU** has continued its research on adaptivity and autonomy in embedded systems. Focus in year 4 has been on three main areas:
    Implementation of fault detection in eDNA prototype

Finishing PhD thesis related to the eDNA project
Preparing eDNA for spin-off company

Part 1 was completed succesfully in the spring of 2011 and resulted in a paper [BKM11], which was submitted and presented at the NASA/ESA Conference on Adaptive Hardware Systems 2011 in San Diego. A paper dealing with optimizing of the placement of spare cells on the eDNA architecture was also presented at the same conference [BMP11]. We have also submitted a journal paper for the Transactions on Aerospace and Electronic Systems, which describe the whole eDNA project [BMK11].

Michael R. Boesen from **DTU** completed his PhD thesis [Boesen11] by May 31st about the eDNA architecture. The PhD thesis was delivered in a 220+ pages monograph form. The thesis was defended in September 2011 and resulted in Michael being awarded the PhD title. Futhermore at the PhD graduation ceremony Michael was awarded the Managing Director P. Gorm Petersen award for outstanding research.

**DTU** is currently trying to commercialize the technology through a spin-off company. During the last 6 months they have succesfully identified a highly scalable market in cooperation with a Danish partner with whom they are currently targeting a launch of an eDNA based product within a 3 year timeframe. **DTU** has also been in contact with several potential investors and are currently working with a subset of these towards an investment within a year. They are also currently working on an application for the Danish Advanced Technology Foundation.

### MPSoC design and programming (DTU)

**DTU** has developed and composed a research multi-core platform. The platform has been shown to outperform proprietary platforms by 29 to 42 percents. The platform is ideally suited for experimentation and will be used for exploring hardware structures.

A port of our own research operating system, FenixOS, to the ARM processor architecture has been developed. We will extend the operating system to support heterogeneous platforms and develop new operating system structures which will reduce the scalability overheads found in current operating systems [SK11, SKb11].

We have continued work on our interactive compiler environment, code comments. The system is based on the open source GCC compiler and allows the programmer to receive feedback from the compiler on how the code can be made more efficient. We have in the last six months enhanced the environment to support feedback on memory performance when using matrixes. The environment has also been revised so that future enhancements can be made more easily [LLKZ11, LKM11].

During year 4, DTU have started a new activity on the mapping of embedded applications onto multi-ASIP platforms in the early stage of the design process. An optimized mapping can be obtained only when the number, interconnection and configuration of ASIPs have been fully determined. However, in order to _find an optimized configuration of the platform and the individual ASIPs, the mapping has to be known. To address this problem, DTU propose a probabilistic estimation technique to capture the uncertainties of the detailed ASIP configurations and to identify the platform configuration and mapping which has the highest probability of meeting the requirements after a detailed ASIP optimization has been applied at a later stage. I.e., addressing the initial macroarchitecture synthesis by abstracting the micro-architecture synthesis by a probability distribution of all possible ASIP configurations. Once the mapping and the platform compositions are defined, this information can be forwarded to a micro-architecture design phase to optimize the ASIP configurations. This exchange of information between the macro and the micro-level can be iterated in order to provide a refinement of both the micro and macro-architectures.

**Bus Access Design for Combined Worst and Average Case Execution Time Optimization of Predictable Real-Time Applications on Multiprocessor Systems-on-Chip (UNIBO, Linköping)**

In real-time systems correctness of a program not only depends on the produced computational results, but also on its ability to deliver these on time, according to specified time constraints.  Software running on these systems grow more and more complex, thus requiring more computational power in terms of hardware resources. In order to satisfy these demands, multi-core systems on a single chip are used to an increasing extent. To achieve predictability with respect to time, various techniques are applied, assuming that the worst-case execution time (WCET) of every task is known, and a lot of research has been carried out within this area The main obstacle when performing timing analysis on multiprocessor systems is that the scheduling of tasks assumes that their worst-case execution times are known. However, to calculate WCET, knowledge about the task schedule is in turn required. Thus, the traditional method of separating WCET analysis and task scheduling no longer works on new architectures, leading to much longer execution times than what can be expected on average. This, in traditional techniques, increases the gap between the worst-case global delay (WCGD) and the average-case global delay (ACGD). In this intervals, non-working processors are left idle and are potentially free to be used for anything, as long as they are ready and free at the start of the new period, or can simply be shut-off, to save energy. Consequently, it is of great interest that the average-case global delay is as short as possible, even for hard real-time systems.

In a joint work, **Linköping** and **UNIBO** introduced a novel off-line methodology for optimizing the average case WCET and ACET simultaneously, allowing for a good average-case execution time while still keeping the worst case as small as possible. An implementation of developed techniques on a MPSoC cycle accurate simulator an exhaustive set of experiments to validate the approach are performed. Tests with real case automotive and avionics application tasks have also demonstrated the effectiveness.

**SoC-TM: Integrated HW/SW Support for Transactional Memory Programming on Embedded MPSoCs**

The shared memory paradigm is widely adopted in embedded MPSoC designs, since it provides an easy-to-understand abstraction of memory resources: a single address space, to which programmers are accustomed. Simplifying thread synchronization in shared memory programs is thus paramount. Fine-grained locking techniques can obtain achieve high performance, but this requires deep understanding of the target application and complicated debug processes. Coarse-grained locking is easier to use, but cannot extract high degrees of parallelism. Transactional Memory (TM) has emerged as a promising approach to address the difficulties of shared memory programming in the general-purpose computing domain. Transactions are as easy to use as coarse-grained locks, but promise the same performance of the finest locking granularity. Transactional threads optimistically execute in parallel, as if no data conflict may possibly arise. If a data conflict takes place, the TM system detects it, and appropriately restarts one of the conflicting transactions ensuring that its changes to memory are not visible to the rest of the system. While transactional memory has been extensively studied for the general-purpose computing domain, there have been relatively few works that consider adopting TM in the embedded domain. Any practical TM design for embedded systems must emphasize simplicity; complex hardware designs that require extensive changes to established protocols (i.e., cache coherency), will likely be too costly to adapt. Similarly, TM programming must be integrated into practical and familiar environments, with simple abstractions, in order to easily use it in an embedded domain. The outcome of the research project is *SoC-TM*, an integrated HW/SW solution for TM on embedded MPSoCs. At the heart of **UNIBO** proposal sits a Hardware Transactional Memory (HTM) design that requires only a

few minor modifications to the cache system, plus a dedicated HW component – the *Bloom module*. This is composed of a collection of Bloom filters and is responsible for keeping track of the history of each transaction and managing conflicts. Ease of use and transparency of the internal functionality are first-class design requirements of the **SoC-TM** system. As such, application developers are not meant to cope directly with low-level transactional programming. Instead, TM features are triggered through the use of a custom set of compiler directives, implemented as an extension to the popular OpenMP programming model. To further improve ease of programming, **UNIBO** framework supports speculative task- and data-level parallelism. Specifically, loops which may carry cross-iteration dependencies can be annotated as parallel loops. The underlying TM system ensures that, in case a real dependence arises, the original sequential program semantics is preserved. This is achieved by means of specific hardware support for committing transactions in program order. Experimental results confirm that **SoC-TM** is a viable and cost-effective solution for embedded MPSoCs, in terms of energy, performance and productivity.

## Power, Thermal and Reliability Aware Resource Management for Multicores Systems (UNIBO, Intel Lab)

High-end multicore processors are characterized by high power density with significant spatial and temporal variability. This leads to power and temperature hot-spots, which may cause non-uniform ageing and accelerated chip failure. These critical issues can be tackled by on-line closed-loop thermal and reliability management policies. Model predictive controllers (MPC) outperform classic feedback controllers since they are capable of minimizing a cost function while enforcing safe working temperature. Unfortunately basic MPC controllers rely on a-priori knowledge of multicore thermal model and their complexity exponentially grows with the number of controlled cores.

To override these limitations **UNIBO** and **Intel** develop  a complete distributed solution that combines energy minimization, MPC based thermal capping and thermal model self-calibration. First, each local node first selects, according with the incoming task workload characteristics, the minimum frequency (fEC) that preserves the performance within a tolerable overhead. Second, a safe working temperature is achieved as the local MPC controller trims the frequency, if needed.  Local controllers jointly optimize global system operation by exchanging a limited amount of information at run-time on a neighbourhood basis. Third we address model uncertainty by self-calibration: each thermal controller node supports learning of the local thermal model by monitoring the thermal response of the neighbourhood area under a set of training stimuli applied during a model calibration phase. The distributed controller strategy combined with the distributed thermal model calibration phase allow us to take advantage of the parallelism of the underlying multi-core platform by running different instances of the controller and self-calibration routine in parallel.

In conjunction with that **UNIBO** and **Intel** designed an advanced energy controller  that performs a temperature aware work load allocation problem over a Multi-Processor Systems on Chip (MPSoC) with Dynamic Voltage and Frequency Scaling (DVFS) capabilities. It can work tightly coupled with the thermal controller.  Indeed DVFS allows the programmer to slow the pace of one or more processors, to let the system cool down and become ready to accept more demanding tasks later on. The thermal behavior of a MPSoC device is the result of the interaction of many concurrent factors (including heat conduction, processor workload, chip layout). Despite the dynamic of the single phenomena is known, the complexity of the overall system makes it very hard to devise a declarative thermal model. In such a context, we design and train a Neural Network to approximate the system thermal behavior. The resulting network has been embedded in a combinatorial model and used to produce an optimized workload allocation, avoiding resource over-heating as well as over-usage. We tested the approach obtaining consistently better result compared to a load balancing strategy guided by a

temperature aware heuristic; moreover, we even improve the results of a very well-performing surrogate temperature measure.

### Run-time monitoring of application parameters (IMEC)

The technique developed by **IMEC** works as follows. In order to trade the output quality for the requirements (e.g., application deadline, power budget of the platform) to obtain the results, a set of well-chosen application parameters are tuned/monitored at run time once the requirements are set by the user. This technique is called iteratively in the application to monitor the application, check whether the requirements are met, and take decisions concerning the tuning of the application parameters. It modifies the application parameters either to improve the output quality if the requirements are met, or to lower the output quality if the requirements are not met. To enable an efficient run-time decision-making, the technique makes use of a statistical model derived at design time from a representative set of available input data. This approach has been demonstrated by IMEC in a Distributed and hierachical run-time resource management for audio-driven surveillance applications on heterogeneous MP-SoC platform: A global run-time manager is being implemented as a software task running in parallel with the application controller on the host processor of the platform. It is used to adapt the platform and application at run time and to find global and optimal trdae-offs in applcation mapping based on a given optimization goal.

*-- The above is new material, not present in the Y3 deliverable --*

## 2.2   Individual Publications Resulting from these Achievements

### TU Braunschweig

[ADNSSE11] Philip Axer, Jonas Diemer, Mircea Negrean, Maurice Sebastian, Simon Schliecker, und Rolf Ernst, "Mastering MPSoCs for Mixed-Critical Applications" *IPSJ Transactions on System LSI Design Methodology*, vol. 4, August 2011

[ASE11] Philip Axer, Maurice Sebastian, und Rolf Ernst, "Reliability Analysis for MPSoCs with Mixed-Critical, Hard Real-Time Constraints" in *Proc. Intl. Conference on Hardware/Software Codesign and System Synthesis (CODES+ISSS)*, (Taiwan), Oktober 2011

[DRNSE11] Jonas Diemer, Jonas Rox, Mircea Negrean, Steffen Stein, und Rolf Ernst, "Real-Time Communication Analysis for Networks with Two-Stage Arbitration" in *Proceedings of the ninth ACM International Conference on Embedded Software (EMSOFT)*, (Taipei, Taiwan), pp. 243-252, ACM, Oktober 2011.

[SAE11] Maurice Sebastian, Philip Axer and Rolf Ernst, "Utilizing Hidden Markov Models for Formal Reliability Analysis of Real-Time Communication Systems with Errors", in *Proceedings of the 17th IEEE Pacific Rim International Symposium on Dependable Computing (PRDC 2011)*, December 2011.

### ETHZ

[TSYB11] Lothar Thiele, Lars Schor, Hoeseok Yang, Iuliana Bacivarov: Thermal-aware system analysis and software synthesis for embedded multi-processors. DAC 2011: 268-273.

[KCT11] Pratyush Kumar, Jian-Jia Chen, Lothar Thiele: Demand bound server: generalized resource reservation for hard  real-time systems. EMSOFT 2011: 233-242

[KT11a] Pratyush Kumar, Lothar Thiele: End-to-End Delay Minimization in Thermally Constrained Distributed Systems. Proc. of 23rd Euromicro Conference on Real-Time Systems (ECRTS 2011), IEEE, Porto, Portugal, July, 2011.

[KT11b] Pratyush Kumar, Lothar Thiele: Cool Shapers: Shaping Real-Time Tasks for Improved Thermal Guarantees. Proc. of Design Automation Conference (DAC 2011), ACM, San Diego, California, June, 2011.

[CHT11] Jian-Jia Chen, Kai Huang, Lothar Thiele: Power Management Schemes for Heterogeneous Clusters under Quality of Service Requirements. 26th Symposium On Applied Computing (SAC), ACM, Taiwan, 2011.


**CEA LIST**

C. Bechara, N. Ventroux and D. Etiemble. Comparison of different thread scheduling strategies for Asymmetric Chip MultiThreading architectures in embedded systems, Euromicro Conference on Digital System Design Architectures, Methods and Tools (DSD), Oulu, Finland, September 2011.

C. Bechara, N. Ventroux and D. Etiemble. AHDAM: an Asymmetric Homogeneous with Dynamic Allocation Manycore chip, Facing the Multicore-Challenge II, Springer-Verlag LNCS 6310, Karlsruhe, Germany, September 2011.

N. Ventroux, T. Sassolas, A. Guerre, B. Creusillet and R. Keryell. SESAM/Par4All: A Tool for Joint Exploration of MPSoC Architectures and Dynamic Dataflow Code Generation, HIPEAC Workshop on Rapid Simulation and Performance Evaluation: Methods and Tools (RAPIDO), Paris, France, January 2012.

J. Peeters, N. Ventroux, T. Sassolas and L. Lacassagne. A SystemC TLM Framework for Distributed Simulation of Complex Systems with Unpredictable Communication, IEEE Conference on Design and Architectures for Signal and Image Processing (DASIP), Tampere, Finland, November 2011.

C. Bechara, N. Ventroux and D. Etiemble. A TLM-based Multithreaded Instruction Set Simulator for MPSoC Simulation Environment, International Workshop on Rapid Simulation and Performance Evaluation: Methods and Tools (RAPIDO), Crete, Greece, January 2011.


**Linköping**

[BAE 2011] Min Bao, Alexandru Andrei, Petru Eles, Zebo Peng, Temperature-Aware Idle Time Distribution for Leakage Energy Optimization, IEEE Transactions on Very Large Scale Integration (VLSI) Systems (accepted for publication).

 [BAE 2011] Min Bao, Alexandru Andrei, Petru Eles, Zebo Peng, On-Line Temperature-Aware Idle Time Distribution for Leakage Energy Optimization, 6th International Symposium on Electronic Design, Test and Applications (DELTA 2011), Queenstown, New Zealand, January, 2011.

[JEP 2011] Ke Jiang, Petru Eles, Zebo Peng, Co-Design Techniques for Distributed Real-Time Embedded Systems with Communication Security Constraints, Design Automation and Test in Europe (DATE12), (accepted for publication).

[JEP 2011] Ke Jiang, Petru Eles, Zebo Peng, Optimization of Message Encryption for Distributed Embedded Systems with Real-Time Constraints, 14th IEEE Symposium on Design and Diagnostics of Electronic Circuits and Systems (DDECS11), Cottbus, Germany, April 13-15, 2011.


**EPFL**

[BCTBPM11] S. Bobba, A. Chakraborthy, O. Thomas, P. Batude, V. Pavlidis, and G. De Micheli. "Performance Analysis of 3-D Monolithic Integrated Circuits". In Proceedings of the IEEE International 3D System Integration Conference (3DIC'11), 2011.

[BCTBEFPM11] S. Bobba, A. Chakraborty, O. Thomas, P. Batude, T. Ernst, O. Faynot, D. Z. Pan, and G. De Micheli. "CELONCEL: Effective Design Technique for 3-D Monolithic Integration targeting High Performance Integrated Circuits". In Proceedings of the 16th Asia and South Pacific Design Automation Conference, pages 337-343, 2011.

**DTU**

[MPM11] E. Maftei, P. Pop, J. Madsen, Module-Based Synthesis of Digital Microfluidic Biochips with Droplet-Aware Operation Execution, accepted for ACM Journal on Emerging Technologies in Computing Systems, 2011.

[SK11] Passas, Stavros; Karlsson, Sven): Comparing the Overhead of Lock-based and Lock-free Implementations of Priority Queues. In: *Proceedings of Forth Workshop on* Programmability *Issues for Heterogeneous Multicores* Presented at: Programmability Issues for Heterogeneous Multicores, Heraklion, Crete

[LLKZ11] Larsen, Per; Ladelsky, Razya; Karlson, Sven; Zaks, Ayal (2011): Compiler Driven Code Comments and Refactoring. In: *Proceedings of Forth Workshop on Programmability Issues for Heterogeneous Multicores* Presented at: Programmability Issues for Heterogeneous Multicores; 4, Heraklion, Crete, Greece

[LKM11] Larsen, Per; Karlsson, Sven; Madsen, Jan (2011): Expressing Coarse-Grain Dependencies Among Tasks in Shared Memory Programs. *I E E E Transactions on Industrial Informatics*, *7*(4), 652-660, doi:10.1109/TII.2011.2166769.

[SK11b] Passas, Stavros; Karlsson, Sven (2011): SRC: FenixOS - A Research Operating System Focused on High Scalability and Reliability. In: *ICS '11: Proceedings of the international conference on Supercomputing* (p. 371) Presented at: International conference on Supercomputing; 25, Tucson, Arizona, USA ACM, doi:10.1145/1995896.1995954.

[TSP11a] Domitian Tamas-Selicean, Paul Pop, "Optimization of Time-Partitions for Mixed-Criticality Real-Time Distributed Embedded Systems", Proceedings of the 14th IEEE International Symposium on Object/Component/Service-Oriented Real-Time Distributed Computing Workshops (ISORCW), 2011.

[TSP11b] Domitian Tamas-Selicean, Paul Pop, "Design Optimization of Mixed-Criticality Real-Time Applications on Cost-Constrained Partitioned Architectures", Proceedings of the IEEE Real-Time Systems Symposium, 2011.

[TSP11c] Domitian Tamas-Selicean, Paul Pop, "Task Mapping and Partitioning Allocation for mIxed-Criticality Real-Time Systems", Proceedings of the IEEE Pacific Rim International Symposium on Dependable Computing (PRDC), 2011.

[MPM11] Wajid Hassan Minhass, Paul Pop, Jan Madsen, "System-level modeling and synthesis of flow-based microfluidic biochips", Proceedings of the CASES conference, p. 225-234, 2011.

[MPMHPD11] W.H. Minhass, P. Pop, J. Madsen, M. Hemmingsen, P. Skafte-Pedersen, M. Dufva. Cell Culture Microfluidic Biochips: Experimental Throughput Maximization. In 5th International Conference on Bioinformatics and Biomedical Engineering, (iCBBE), 2011

[BKM11] Michael Reibel Boesen, Didier Keymeulen, Jan Madsen, Thomas Lu and Tien-Hsin Chao. *Integration of the self-healing eDNA architecture in a liquid crystal waveguide-based fourier transform spectrometer*. NASA/ESA Conference on Adaptive Hardware Systems 2011.

[BMP11] Michael Reibel Boesen, Jan Madsen and Paul Pop. *Application-aware optimization if redundant ressources for the reconfigurable self-healing eDNA hardware architecture.* NASA/ESA Conference on Adaptive Hardware Systems 2011.

[BMK11] Michael Reibel Boesen, Jan Madsen and Didier Keymeulen. *A reconfigurable hardware architecture for high-reliability applications.* Submitted to Transactions on Aerospace and Electronic Systems.

[Boesen11] Michael Reibel Boesen. *A bio-inspired self-healing reconfigurable hardware architecture - Concept, design, prototype and evaluation.* PhD Thesis.

[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

[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.

[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

## UNIBO

Paterna F., Acquaviva A., Caprara A., Papariello F., Desoli G., Benini L., "Variability-Aware Task Allocation for Energy-Efficient Quality of Service Provisioning in Embedded Streaming Multimedia Applications",. ; in:  Computers, IEEE Transactions on, Volume :  PP ,  Issue:99, On page(s): 1, ISSN :  0018-9340. IEEE Press 2011 – DOI: 10.1109/TC.2011.127

Paterna F., Acquaviva A., Caprara A., Papariello F., Desoli G., Benini L., "An Efficient On-line Task Allocation Algorithm for QoS and Energy Efficiency in Multicore Multimedia Platforms" in Design, Automation and Test in Europe Conference and Exhibition, IEEE Proceedings of the Conference on, Grenoble, France, Mar. 15-17, 2011, pp: 100-105.

Bartolini A. ; Cacciari M. ; Tilli A. ; Benini L. ; A distributed and self-calibrating model-predictive controller for energy and thermal management of high-performance multicores, This paper appears in:  Design, Automation & Test in Europe Conference & Exhibition (DATE), 2011, Issue Date :  14-18 March 2011, On page(s): 1 – 6, ISSN :  1530-*1591, Print ISBN: 978-1-61284-208-0. IEEE Press* 2011

Beneventi F.;  Bartolini A.;   Benini L.; Static Thermal Model Learning for High-Performance Multicore Servers,  in: Computer Communications and Networks (ICCCN), 2011 Proceedings of 20th International Conference on, Issue Date: July 31 2011-Aug. 4 2011, On page(s): 1 – 6, Location: Lahaina, HI, USA, ISSN: 1095-2055  Print ISBN: 978-1-4577-0637-*0, 2011 IEEE*

---

*-- The above are new references, not present in the Y3 deliverable --*

---

## 2.3 Interaction and Building Excellence between Partners

**Synthesis Tool for Analysis of 3-D NoC systems (EPFL, UNIBO, STM, ETHZ)**
The interaction in this activity has been between EPFL and UNIBO, working together towards on the development of a synthesis tool for 3-D NoC systems
- There has been active collaboration during the development of the tool.
- Tool development to support power management solutions to design low-power 2-D and 3-D NoC based systems [SMBM11].
- The integration in the tool of an analytical model to design best-effort NoC-based systems with hard and real-time latency constraints [SRMBM11].
- Meetings and visits between UNIBO and EPFL have been made to determine the specifics of the tool in order to improve the access to DRAM memories by means of tuning the system interconnect [SMBM11b].

EPFL and UNIBO are actively participating to the Pro3D project together with STM and ETHZ which focuses on building software tools for next generation 3D-integrated SoC platforms.

**IMEC, ICCS, and STMicroelectronics** are partners of the 2PARMA consortium. The 2PARMA project aims at overcoming the lack of parallel programming models and run-time resource management techniques to exploit the features of manycore processor architectures.

**IMEC and STMicroelectronics are partners of the COMPLEX consortium**
The main objective of the COMPLEX project is to increase the competitiveness of the European semiconductor, system integrator and EDA industry by addressing the problem of platform-based design space exploration under consideration of power and performance constraints early in the design process. High performance usually causes high power consumption. A main challenge in today's embedded system design is to find the perfect balance between performance and power. This balance can not be found efficiently and at high quality, because until now no generic framework for accurately and jointly estimating performance and power consumption starting at the algorithmic level is available. This can only be achieved in cooperation on a European level, taking into account European platform providers, system developers/integrators, EDA companies, Universities and research institutes from both, the HW and the embedded SW world.

---

*-- Changes wrt Y3 deliverable --*

*This section has been focused on NEW interaction being built in the scope of newly funded EU projects, and highlighting the active participation of industry in these activities.*

---

## 2.4  *Joint Publications Resulting from these Achievements*

### TUBS, Symtavision

[SAEFJ11] Maurice Sebastian, Philip Axer, Rolf Ernst, Nico Feiertag, und Marek Jersak, **"Efficient Reliability and Safety Analysis for Mixed-Criticality Embedded Systems"** in SAE 2011 World Congress & Exhibition Technical Paper, Detroit, USA, April 2011.

### ETHZ, TU Dortmund

[MTKBTHLXH11] Peter Marwedel, Jürgen Teich, Georgia Kouveli, Iuliana Bacivarov, Lothar Thiele, Soonhoi Ha, Chanhee Lee, Qiang Xu, Lin Huang: Mapping of applications to MPSoCs. CODES+ISSS 2011: 109-118.

### ETH, SSSA PISA

Kai Huang, Luca Santinelli, Jian-Jia Chen, Lothar Thiele, Giorgio C. Buttazzo: Applying real-time interface and calculus for dynamic power management in hard real-time systems. Real-Time Systems, Springer Netherlands, Vol. 47, No. 2, pages 163-193, March, 2011.

### EPFL, UNIBO

[SMBM11] C. Seiculescu, S. Murali, L. Benini, and G. De Micheli. "Design and Analysis of NoCs for Low-Power 2D and 3D SoCs". In C. Silvano, M. Lajolo, and G. Palermo, editors, Low Power Networks-on-Chip, number 3, pages 199-222, Springer New York, 2011.

[SRMBM11] C. Seiculescu, D. Rahmati,  S. Murali, L. Benini, and G. De Micheli, H. Sarbazi-Azad. "Designing Best Effort Networks-on-Chip to Meet Hard Latency Constraints", Accepted to be published in ACM Transactions on Embedded Computer Systems.

[SMBM11b] C. Seiculescu, S. Murali, L. Benini, and G. De Micheli, "A DRAM Centric NoC Architecture and Topology Design Approach". In Proceedings of the IEEE Computer Society Annual Symposium on VLSI, pages 54-59, 2011.

### UNIBO, Linköping

Jakob Ròsen, Carl-Fredrik Neikter, Petru Eles, Zebo Peng, Paolo Burgio and Luca Benini , "Bus Access Design for Combined Worst and Average Case Execution Time Optimization of Predictable Real-Time Applications on Multiprocessor Systems-on-Chip", in: Real-Time and Embedded Technology and Applications Symposium (RTAS), 2011 17th IEEE, Chicago, IL, pages 291 - 301, 2011

### IMEC, STM-I

C. Silvano, W. Fornaciari, S. Crespi Reghizzi, G. Agosta, G. Palermo, V. Zaccaria, P. Bellasi, F. Castro, S. Corbetta, A. Di Biagio, E. Speziale, M. Tartara, D. Melpignano, J.-M. Zins, D. Siorpaes, H. Hubert, B. Stabernack, J. Brandenburg, M. Palkovic, P. Raghavan, C. Ykman-Couvreur, A. Bartzas, S. Xydis, D. Soudris, T. Kempf, G. Ascheid, R. Leupers, H. Meyr, J. Ansari, P. Mahonen and B. Vanthournout, "2PARMA: Parallel Paradigms and 3 Run-time Management Techniques for Many-Core Architectures", pp. 65-79, in "VLSI 2010 Annual Symposium", Selected Papers, Editors: N. Voros, A. Mukherjee, N. Sklavos, K. Masselos, M. Huebner, Lecture Notes in Electrical Engineering, Volume 57, 1st Edition., 2011, VIII, 331 p., Springer Netherlands, August 31, 2011, ISBN 978-94-007-1487-8

C. Silvano, W. Fornaciari, S. Crespi Reghizzi, G. Agosta, G. Palermo, V. Zaccaria, P. Bellasi, F. Castro, S. Corbetta, E. Speziale, D. Melpignano, J.M. Zins, H. Hubert, B. Stabernack, J. Brandenburg, M. Palkovic, P. Raghavan, C. Ykman-Couvreur, A. Bartzas, D. Soudris, T. Kempf, G. Ascheid, H. Meyr, and J. Ansari, P. Mahonen, B. Vanthournout, "Parallel Paradigms and Run-time Management Techniques for Many-core Architectures: The 2PARMA Approach",

accepted for presentation in IEEE 9th International Conference on Industrial Informatics (INDIN'2011), Caparica, Lisbon, Portugal, July 26-29 2011 *(Invited Paper)*.

A. Bartzas, P. Bellasi, I. Anagnostopoulos, C. Silvano, W. Fornaciari, D. Soudris, D. Melpignano, C. Ykman-Couvreur, "Runtime Resource Management Techniques for Many-core Architectures: The 2PARMA Approach", *accepted* for presentation in the International Conference on Engineering of Reconfigurable Systems and Algorithms (ERSA'11), Monte Carlo Resort, Las Vegas, Nevada, USA, July 18-21 2011 *(Invited Paper)*.

C. Silvano, W. Fornaciari, S. Crespi Reghizzi, G. Agosta, G. Palermo, V. Zaccaria, P. Bellasi, F. Castro, S. Corbetta, E. Speziale, D. Melpignano, J.M. Zins, H. Hubert, B. Stabernack, J. Brandenburg, M. Palkovic, P. Raghavan, C. Ykman-Couvreur, I. Anagnostopoulos, A. Bartzas, D. Soudris, T. Kempf, G. Ascheid, H. Meyr, J. Ansari, P. Mahonen, B. Vanthournout, "Parallel programming and Run-time Resource Management Framework for Many-core Platforms: The 2PARMA Approach", accepted for presentation in the 6th International Workshop on Reconfigurable Communication-centric Systems-on-Chip (ReCoSoC'2011), Montpellier, France, June 20-22 2011 (invited paper).

Fabien Colas-Bigey (Thales), Sara Bocchio (ST-I), Chantal Couvreur (IMEC), Gianluca Palermo (POLIMI),Philipp A. Hartmann (OFFIS) "Audio Driven Video Surveillance System using COMPLEX design flow", DATE workshop on designing for embedded parallel computing platforms: architectures, design tools, and applications, Grenoble, March 2011.

*-- The above are new references, not present in the Y3 deliverable --*

## 2.5   Keynotes, Workshops, Tutorials

**Invited Talk: Trusted MpSoC Platforms for Safety Related Applications**
**(Rolf Ernst, TU Braunschweig)**
**Coolchips Symposium**
*Yokohama, Japan – April 20 - 22, 2011(per video conference)*

COOL Chips is an International Symposium initiated in 1998 to present advancement of low-power and high-speed chips. The symposium covers leading-edge technologies in all areas of microprocessors and their applications.

MpSoCs are efficient platforms for systems integration. However, due to physical resource sharing, safety critical systems integration becomes more challenging compared to distributed systems potentially leading to increasing design and certification cost. There are many issues in efficient error detection and handling, function segregation, timing, or in the RTE interface that require system level hardware/software solutions. Integration frequently leads to mixed critical systems, i.e. systems which combine functions of different criticality levels. One of the main requirements is a flexible trade-off between development, certification, and production cost that is highly influenced by production volume, product lifetime, and system criticality. The talk gave an introduction to the design of safety critical systems, and the derived design challenges for MpSoC based systems. Novel solutions and tool technologies were explained. The talk also gave an overview on several large industrial-academic projects that deal with safety critical MpSoC and their application, such as the ARTEMIS RECOMP project (Reduced Certification Costs for Trusted Multi-core Platforms).

http://www.coolchips.org/

**Keynote: Embedded Systems - the Neural Backbone of Society**
**(Rolf Ernst, TU Braunschweig)**
**ARTEMIS Strategic Research Agenda Symposium**
*Brussels, Belgium – May 18, 2011*

The technology of Embedded Systems is adding intelligence to all kind of objects. Due to the openness of the Internet, Embedded Systems will also be able to "get access to information systems and in turn the information systems get access to the Embedded Systems which now enables the internet of things". With communication, Embedded Systems have gained a strategic role and networked Embedded Systems are now considered the neural system of the digital society. Keynote speaker Prof. Rolf Ernst from the TU Braunschweig spoke about embedded systems as the neural backbone of society.

http://www.artemis-ia.eu/sra_home

**MiniKeynote: MpSoC for safety critical applications – from multicore to manycore**
**(Rolf Ernst, TU Braunschweig)**
**11th EDAA/IEEE Forum on Embedded MPSoC and Multicore**
*Beaune, France – July 8, 2011*

In 2010 Prof. Ernst gave a short introduction on requirements and design methods for MpSoC in safety critical applications. The focus was on interference of safety critical and non-critical applications via shared resources and the corresponding requirements imposed by safety standards. In many-core systems interference is even stronger due to multi-hop NoCs and memory hierarchies. The talk in 2011 gave an overview on first results of a research platform under development as part of the European ARTEMIS project RECOMP.

http://www.mpsoc-forum.org/previous/2011/agenda.html

**Invited Talk: Formal Performance Analysis in Automotive Systems Design – A Rocky Ride to New Grounds**
**(Rolf Ernst, TU Braunschweig)**
**23rd IEEE Conference on Computed Aided Verification (CAV) Symposium**
*Snowbird, Utah, USA – July 20, 2011*

CAV 2011 was the 23rd in a series dedicated to the advancement of the theory and practice of computer-aided formal analysis methods for hardware and software systems. The conference covered the spectrum from theoretical results to concrete applications, with an emphasis on practical verification tools and the algorithms and techniques that are needed for their implementation. The talk given by Prof. Ernst focused on performance challenges in automotive design. Formal performance analysis methods for automotive design were presented and major obstacles from theory to industrial application were highlighted.
http://www.cs.utah.edu/events/conferences/cav2011/

**Presentation: IDAMC NoC – Efficient Quality-of-Service Support for Mixed-Critical Networks-on-Chip**
**RECOMP Technical Day**
*Porto, Portugal - August 29, 2011*.

Jonas Diemer (TU Braunschweig) gave a presentation on the Network-on-Chip used in the IDAMC platform.

**Panel Session: ARTEMIS, from successful R&D to cutting-edge Innovation**
**(Rolf Ernst, TU Braunschweig)**
**ARTEMIS Strategic Research Agenda Symposium**
*Brussels, Belgium – October 4, 2011*

In the panel session held at the European Parliament Rolf Ernst talk about the key role of embedded systems in industrial innovation. Furthermore it highlighted the benefits of ARTEMIS from the previous research projects on embedded systems e.g. FP6 and FP7.

http://www.artemis-ia.eu/jti_programme

**Invited Talks:**
  1. **Timing Analysis of Ethernet AVB for Real-Time Systems (Jonas Diemer)**
  2. **Combining Security with Reliability using Multi-core (Philip Axer)**
**Symtavision News Conference**
*Braunschweig, Germany - October 5, 2011*

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. This year continued with the successful implementation of a technical day with parallel practice and research tracks. TU Braunschweig was invited to present current research results on real-time analysis methods for multi-core systems.

Jonas Diemer (TU Braunschweig) gave a presentation on the timing analysis of Ethernet AVB. Philip Axer (TU Braunschweig) gave a presentation on the challenges of reliability and security in multicore systems for safety-critical applications

**Invited Talk: Resilient Real Time OS**

**(Rolf Ernst, TU Braunschweig)**
*CODES+ISSS 2011*
*Taipei, Taiwan – October 9 - 14, 2011*

Prof. Ernst gave a talk in the Special Session Design and Architecture for Dependable Embedded Systems of the Embedded System Week event. As Moore's Law advances to sub-50nm technology nodes, reliability of integrated circuits becomes an inherent problem: effects like electro-migration, SEUs (Single-Event Upsets), thermal effects due to increased power density and others all have a negative impact on reliability either in form of transient or permanent faults or degrading system characteristics (aging effects). The reliability problem can be addressed at various levels of abstraction. The special session aimed at addressing the interdependency between the various levels of abstraction.

http://www.esweek.org/

**Invited Talk: Multicore Architectures for Mixed Safety Critical Applications – Challenges and Opportunities**
**(Rolf Ernst, TU Braunschweig)**
**SafeTRANS Industrial Day**
*Hamburg, Germany – November 08, 2011*

**SafeTRANS** ("Safety in Transportation Systems") is a Competence Cluster combining research and development expertise in the area of complex embedded systems in

transportation systems. SafeTRANS drives research in human centred design, in system and software development methods for embedded systems, as well as in safety analysis and - for avionics and rail - its integration in certification processes, driven by a harmonised strategy addressing the need of the transportation sector. The topic of the 11[th] SafeTRANS Industrial Day was "Development processes for Multicore".

Sharing embedded system resources among functions of different safety criticality usually leads to mixed safety and time critical embedded systems. Such mixed critical systems must combine conflicting safety and efficiency requirements and related design processes. The talk gave an overview on mixed critical system design challenges and explained how different criticalities can be properly separated in function and timing. In multicore architectures, separation is more difficult than in networks due to low level resource sharing. The talk showed the effects and provided solutions addressing multicore and manycore systems.

http://www.safetrans-de.org/en_11_Industrial_Day.php

**Tutorial / Invited Talk: Multi-Core and Many-Core for Mixed-Critical Systems - Denial of Service and other Performance Challenges**
**(Mircea Negrean, Rolf Ernst, TU Braunschweig)**
**BoCSE (Bosch Conference on Systems and Software Engineering)**
*Ludwigsburg, Germany – November 15 – 17, 2011*

The talk was part of a tutorial at the 4[th] BoCSE-Conference. The conference organized by Bosch brings toghether engineers, managers, technology experts from different departments of the company, from other companies and from academia. In 2011 the event had over 600 participants. The focus of the given presentation was on challenges which arise in case of integrating applications with different criticalities/different safety requirements on multi-core and many-core systems.

**Invited Talk: The mixed criticality challenge to embedded system platforms**
**(Rolf Ernst, TU Braunschweig)**
**ICT.OPEN**
*Veldhoven, Netherlands – November 14 - 15, 2011*

ICT.OPEN is the principal ICT and Computer Science research conference in the Netherlands. It features plenary key notes and invited speakers, as well as selected oral and poster presentations. The state of art in ICT and Computer Science research is presented and discussed and therefore ICT.OPEN aims to be the place to be for everybody involved or interested in ICT and Computer Science research.

The talk discussed challenges in the design of mixed critical systems with a focus on multi-core architecture. First solutions proposed in major projects which address the mixed-criticality challenge in the larger context of automotive electronics and smart buildings have been presented.

http://www.nwo.nl/nwohome.nsf/pages/NWOP_8M3AYV

**Invited Talk: Synthesis and optimization in mixed critical systems**
**(Rolf Ernst, TU Braunschweig)**
**SOMRES Workshop 2011**
*Vienna, Austria – November 29, 2011*

The Workshop on Synthesis and Optimization Methods for Real-time Embedded Systems (SOMRES) was part of the 32nd IEEE Real-Time Systems Symposium (RTSS). The event was

214373 ArtistDesign NoE   JPRA   Year 4 (Jan-Dec 2011)
Cluster: Hardware Platforms and MPSoCs  D12-(6.1)-Y4
Activity: Platform and MPSoC Design

organized with Artist partners. Prof. Ernst gave a talk on synthesis and optimization in mixed critical systems.

http://www.artist-embedded.org/artist/SOMRES-20111.html and http://retis.sssup.it/synthesys/


## Lothar Thiele, Iuliana Bacivarov: Thermal Aware Mapping for MPSoCs

DAC Workshop on Multiprocessor System-on-Chip for Cyber Physical Systems: Programmability, Run-Time Support, and Hardware Platforms for High Performance Embedded Applications.

*DAC 2011. Thursday, June 9, 2011, Time: 8:30 AM — 5:00 PM*

Cyber physical systems combine the system's computational with the physical elements. This type of systems can often be found in embedded systems of various domains like aerospace, automotive, chemical processes, civil infrastructure, energy, healthcare, manufacturing, transportation, entertainment, and consumer appliances. Since applications in this domains require a high computational performance, novel approaches for the computing elements have to be provided. Multicore hardware is a promising solution to provide the sufficient performance / power consumption ratio. Multicore architectures offer a better performance/Watt ratio than single core architectures with similar performance. Combining multicore and coprocessor technology promise extreme computing power for highly CPU-time-consuming applications. Especially, FPGA-based accelerators not only offer the opportunity to speedup an application by implementing their compute-intensive kernels into hardware but also to adapt to the dynamical behavior of an application. The purpose of the third edition of this very successful workshop is to evaluate strategies for future system design in MPSoC architectures, especially for cyber physical systems. Both aspects, hardware design and tool-integration into existing development tools will be discussed. Especially, the novel trends in MPSoC combined with reconfigurable architectures are a topic in this workshop. The main emphasis is on architectures, design-flow, tool-development, applications and system design.


## Invited Talk: Lothar Thiele ETH Zurich: Temperature-aware Scheduling

ARTIST Summer School September 4-9, 2011

*Aix-les-Bains (near Grenoble), France*

Power density has been continuously increasing in modern processors, leading to high on-chip temperatures. A system could fail if the operating temperature exceeds a certain threshold, leading to low reliability and even chip burnout. There have been many results in recent years about thermal management, including (1) thermal-constrained scheduling to maximize performance or determine the schedulability of real-time systems under given temperature constraints, (2) peak temperature reduction to meet performance constraints,and (3) thermal control by applying control theory for system adaption. The presentation will cover challenges, problems and approaches to real-time scheduling under temperature constraints for single- as well as multi-processors.


## Mini-keynote: "Hardware support for online resources management " (Raphaël David, CEA LIST) + Program Co-Chair
## International Forum on Embedded MPSoC and Multicore, MPSoC'2011
*Beaune, France, July 4-8, 2011*

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 2011, Raphaël David serves as program chair and in his talk he presented some of the main challenges regarding dynamic

management of parallel systems and discussed the Hardware Synchronizer resource and its usefulness for advance MPSoC.

http://www.mpsoc-forum.org/previous/2011/program.html

**Tutorial: "Dynamic management of Embedded Multi-core architectures" (Raphaël David, CEA LIST)**
**Asia South Pacific Design Automation Conference, ASP-DAC'2011**
*Yokohama, Japon, January 25-28, 2011*

Tutorial dedicated to present solutions for dynamically managing computing resources in MPSoC architectures as well as hardware supports for accelerating this management. Focus on CEA LIST experience (SCMP, Platform 2012 architectures and the related acceleration modules).

**Invited Talk, Petru Eles,** Scheduling and Optimization of Fault-Tolerant Embedded Systems
**ACM SIGPLAN/SIGBED Conference on Languages, Compilers, Tools and Theory for Embedded Systems (LCTES 2011), Chicago, IL, USA, 12-14 April 2011 (in conjuction with CPS Week 2011)**

Abstract

This work addresses the issue of design optimization for fault-tolerant hard real-time systems. In particular, our focus is on the handling of transient faults using both checkpointing with rollback recovery and active replication. Fault tolerant schedules are generated based on a conditional process graph representation. The formulated system synthesis approaches decide the assignment of fault-tolerance policies to processes, the optimal placement of checkpoints and the mapping of processes to processors, such that multiple transient faults are tolerated and the timing constraints of the application are satisfied. We propose a fine-grained transparent recovery, where the property of transparency can be selectively applied to processes and messages. Transparency hides the recovery actions in a selected part of the application so that they do not affect the schedule of other processes and messages. While leading to longer schedules, transparent recovery has the advantage of both improved debuggability and less memory needed to store the fault-tolerant schedules.

**Keynote: De Micheli, Giovanni – Nanosystems: devices, circuits, architectures and applications.**
**International SoC Design Conference (ISOCC)**
Jeju, Korea, November 17 – 18, 2012
Abstract: Much of our economy and way of living will be affected by nanotechnologies in the decade to come and beyond. Mastering materials at the molecular level and their interaction with living matter opens up unforeseeable horizons. This talk deals with how we will conceive, design and use nanosystems, i.e., integrated systems exploiting nanodevices. Whereas switching circuits and microelectronics have been the enablers of computer and communication systems, nanosystems have the potentials to realize innovative computational fabrics whose applications require broader hardware abstractions, extended software layers and with a much higher complexity level overall. The abstraction of computation, the nanosystem architecture, the technological feasibility envelope and the multivariate design optimization problems pose challenging and disruptive research questions that this talk will address.
http://www.isocc.org/keynote/sub_04_1.asp

**Tutorial: Model-based MPSoC Architecture Synthesis for Highly-demanding Embedded Applications (Jan Madsen, DTU, Menno Lindwer, Silicon Hive (Intel), Lech Jozwiak, TUe)**
*Grenoble, France - march. 2011*

This tutorial focuses on mastering the automatic architecture synthesis and application mapping for heterogeneous customizable multi-processor systems-on-chip (MPSoCs) based on configurable and extensible application-specific instruction-set processors (ASIPs). The tutorial presents the results of our analysis of the main problems that have to be solved and challenges to be faced in design of such heterogeneous customizable MPSoCs for modern demanding applications. In particular, it discusses the problems of multi-objective optimal architecture synthesis and application mapping, adequate exploitation of multiple design trade-offs, and coherent development of computing, communication and memory sub-systems for complex real-time MPSoCs. It proposes the model-based semi-automatic architecture synthesis methods and EDA-tools that enable effective and efficient solution of these problems.

## Summer school: Platform-Based Design: From Multi-Core Platforms to Biochips and beyond (Jan Madsen, DTU)
**Artist Summer School in China**
*Beijing, China, August 8-12, 2011*
One of the challenges in modern embedded system design is to map the application onto a platform such that essential requirements are met. In order to do so at an early stage in the design process, where not all parts have been implemented or even designed, a system-level model of the application executing on the platform is needed. This model should allow for an accurate modeling of the global performance of the system, including the interrelationships among the diverse processing elements, physical interfaces and inter-connections. This course gives an introduction to the problem of mapping applications onto platforms and how it can be extended to the design of complex biochips.

## Tutorial: Digital Microfluidic Biochips: Functional Diversity, More than Moore, and Cyberphysical Systems (Krishnendu Chakrabarty, Duke University, USA; Paul Pop, DTU; Tsung-Yi Ho, National Cheng Kung University Tainan, Taiwan)
*Taipei, Taiwan - 9.11.2011*
Advances in droplet-based "digital" microfluidics have led to the emergence of biochip devices for automating laboratory procedures in biochemistry and molecular biology. These devices enable the precise control of nanoliter-volume droplets of biochemical samples and reagents. As a result, non-traditional biomedical applications and markets (e.g., high-throughout DNA sequencing, portable and point-of-care clinical diagnostics, protein crystallization for drug discovery), and fundamentally new uses are opening up for ICs and systems. However, continued growth (and larger revenues resulting from technology adoption by pharmaceutical and healthcare companies) depends on advances in chip integration and design-automation tools.
http://www2.imm.dtu.dk/~pop/codes+isss02tu-chakrabarty.html

## Keynote : A Time-predictable Micoprocessor: the Patmos Approach (Martin Schoeberl, DTU)
**11^th International Forum on Embedded MPSoC and Multicore**
*Beaune, France, july 4-8, 2011*
Current processors are optimized for average case performance, often leading to a high worst-case execution time (WCET). Many architectural features that increase the average case performance are hard to be modeled for the WCET analysis. We present Patmos, a processor optimized for low WCET bounds rather than high average case performance. Patmos is a dual-issue, statically scheduled RISC processor. The instruction cache is organized as a method cache and the data cache is organized as a split cache in order to simplify the cache WCET analysis. To fill the dual-issue pipeline with enough useful instructions, Patmos relies on a customized compiler. The compiler also plays a central role in optimizing the application for the

WCET instead of average case performan.

http://www.mpsoc-forum.org/previous/2011/index.html

## Invited Talk: The Self-Healing Computer (Jan Madsen, DTU)
## Seminar in the Danish Engineering Association
*Copenhagen, Denmark, September 14, 2011*

The basic concepts and implementation of the eDNA technology was presented to an audience of engineers from the Danish industry.

## Invited Talk: Recent Research and Emerging Challenges in the System-Level Design of Digital Microfluidic Biochips (*Paul Pop,  Elena Maftei,  Jan Madsen, DTU*)
*SOCC 2011, Taipei, Taiwan*

"Digital" biochips are manipulating liquids as droplets on an array of electrodes. So far, researchers have assumed that microfluidic operations are executing on modules, formed by grouping adjacent electrodes. However, operations can execute by routing droplets on any sequence of electrodes. This paper presents recent work on digital biochip synthesis.

http://www.ieee-socc.org/SOCC2011/Program/program.html

## DATE workshop on designing for embedded parallel computing platforms: architectures, design tools, and applications, Grenoble, March 2011.

*Organizers: Christina Silvano (Polimi), M. Palesi (Kore University), Ch. Ykman-Couvreur (IMEC), D. Gohringer (Fraunhofer IOSB), J. Becker and M. Hubner (Karlsruhe inst. Of Technology)*

Embedded computing is shifting to multi/many-core designs to boost performance due to unacceptable power consumption and operating temperature increase of fast single-core CPU's. Hence, embedded system designers are increasingly faced with several big challenges, namely: the support for a variety of concurrent applications, and the platform heterogeneity. These challenges lead to the following significant design issues: how can applications that exploit the underlying (parallel) architecture be written without burdening the application designer? What does the application designer really need to know of the underlying architecture? What tools are needed to efficiently map applications and what part of the mapping process should/could be automated? How should we design and optimize the underlying architectures? This workshop brings together researchers and practitioners actively working on architectures, design tools, and applications for embedded parallel computing platforms to address these questions and related issues.

**Keynote: Many-core Interconnection Networks Trends: Fast, Vertical, Asynchronous (Luca Benini, UNIBO)**

**International Workshop on System-Level Intercnnect Prediction 2011.**

**http://www.sliponline.org/**

**Invited talk: Going up: 3D integration and many-core SoCs (Luca Benini, UNIBO)**

**3D Integration Workshop For High Performance Computing Systems**

**http://eeweb.poly.edu/hli/3D-Workshop/Home.htm**

# 3.  Milestones, and Future Evolution

## 3.1  Problems to be Tackled in the future

**System-wide reliability analysis with mixed-criticality workloads (TU Braunschweig)**

In the future TU Braunschweig will continue the work on the reliability analysis on system as well as on component layer. From industrial side the need for reliability analysis of gateways has already been indicated. It is the goal to transfer the established methods introduced for reliability analysis of mult-core CPUs to the gateway domain.

The CERTAINTY project should produce first results concerning integration of all these individual analysis steps into an overall system design methodology. Especially the definition of hardware and software requirements should be finalized. Additionally is it planned to specify the final CERTAINTY system model, reflecting all properties relevant for the upcoming work in CERTAINTY including error models or isolation mechanisms for example.

**Integrated Dependable Architecture for Many Cores (TU Braunschweig)**

TU Braunschweig will continue its work on the IDAMC platform for the RECOMP project. We will enhance the platform with new mechanisms for isolation, virtualization and monitoring and bring up a corresponding software infrastructure. This effort is also used in the ASTEROID project which focuses on the reliability aspect. Furthermore, the BMBF-funded ARAMIS project will soon start which will also use the IDAMC infrastructure.

**CEA LIST** will continue investigating solutions for dynamically manage computing and memory resources. Future evolution of **CEA LIST** researchs will first consist in enhancing the features of its Hardware Synchronizer (HWS). A deep evaluation of its performance related to the acceleration of runtime software services will allow to assess its usefuleness in future manycore architectures. Scalability of this approach based on hierarchical control strategies will be studied in the scope of the **COBRA** project. Investigations on hardware accelerators for preemption and migration at the FC (Fabric Controller) level coupled to hardware acceleration at cluster level thanks to the HWS will be at the heart of **CEA LIST** contribution at the design level to **COBRA** project during next year. 2012 will also be the occasion to experiment with the first physical ASIC implementation of a **Platform 2012** multi-cluster implementation (designed as a multi-chip device). This implementation is expected to be available before mid-2012. Finally, a great achievement for **CEA LIST** will be the complete coupling between its precise and high-performance SESAM virtual prototyping platform and commercial emulation solutions, allowing the development and evalution of mixed hardware/virtual implementations of multi/manycore architectures

**Run-time monitoring of application parameters (IMEC)**

Within the European 2PARMA project, the developed technique will be validated on a multiview video processing application mapped on several embedded architectures. Robustness mechanism will be developed too on top of this technique.

**Distributed and hierachical run-time resource management for audio-driven video surveillance applications on heterogeneous MP-SoC platform (IMEC)**

Within the European COMPLEX project, the integration of the gloabl run-time manager in the audio-driven video surveillance application mapped on a heterogeneous MP-SoC platform will be continued.

**Develop monitoring and control techniques in the middleware of the system to automatically adapt platform services to application requirements (IMEC).**

Within the European project Pharaon, the basic blocks will be developed for efficient run-time reconfiguration of multiple standard applications, such software defdined radios.

---

*-- Changes wrt Y3 deliverable --*

*For several partners, the main objective area to tackle beyond ArtistDesign are the same of Y4. 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**

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.

*This milestone has been achieved. Especially the research on system reliability analysis has been intensified by means of the CERTAINTY project. The applicability of analysis for practical problems has been demonstrated based on the industrial-driven example systems provided by Toyota-ITC.*

▪  **Predictability, power, temperature aware MPSoC arcitectures and programming models**

Several steps forward have been made in this milestone.  The foucs of year 4 has been on the mapping and dynamic applications, as well as managing memory bottlenecks. The programming model links with the power and especially temperature aspects have not yet been addressed and are scope for future research*. Industry involvement (e.g. STmicroelectronics and Intel shows increasing interest from companies in the exploitation of the results coming from the ARTIST-DESIGN research community.*

### 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 4  this activity has benefited from direct funding from:

**TU Braunschweig**

- **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. Mixed-criticality systems can be found in a multitude of different areas of application. http://www.ida.ing.tu-bs.de/en/research/projects/recomp/

- **Autonomous Integrated Systems (AIS) Project**

Within the German research project "Autonomous Integrated Systems" (AIS) new design methodologies are explored to tackle the challenges resulting from unreliable components. The project is funded half by the German "Federal Ministry of Education and Research and half by an industry consortium arranged within the "edaCentrum".
http://www.edacentrum.de/ais

- **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 projects have been starten in 2009, 2010 and 2011.

- **CERTAINTY**

The goal of the FP7 CERTAINTY is to develop new system design and analysis methods to support certification procedures for mixed-criticality systems. CERTAINTY addresses different aspects such as static analysis methods for execution time or interference estimation, inclusion of probabilistic information to model and quantify uncertainties such as transient or permanent errors as well as optimized synthesis and deployment strategies. The overall methodology should be composable and scalable to embedded systems of arbitrary size. From the industrial side CERTAINTY is driven by requirements from the aerospace industry.

**CEA LIST**
- **MC2H (Many-core for Computing and Healing). French R&D cooperation program (Nano 2012)**

- **SCALOPES (SCAlable LOw Power Embedded platformS). ARTEMIS project**
- **COBRA (Computing Fabric For High Performance Applications). CATRENE project**
- **SMECY (Smart Multicore Embedded Systems). ARTEMIS project**

**DTU**

- **DaNES** (Danish Network for Embedded Systems, funded by the Danish Advanced Technology Foundation), Denmark. Period 2007-2011.
- **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.

**IMEC**

### 2PARMA project

The number of cores to be integrated in a single chip is expected to increase rapidly in the coming years, moving from multi-core to many-core architectures. This trend will require a global rethinking of software and hardware design approaches. This class of computing systems (Manycore Computing Fabric) promises to increase performance, scalability and flexibility if appropriate design and programming methodologies will be defined to exploit the high degree of parallelism exposed by the architecture. Other potential benefits of Many-core Computing Fabric include energy efficiency, improved silicon yield, and accounting for local process variations. To exploit these potential benefits, effective run-time power and resource management techniques are needed. With respect to

conventional computing architectures, Many-core Computing Fabric offers some customisation capabilities to extend and/or configure at run-time the architectural template to address a variable workload. The 2PARMA project aims at overcoming the lack of parallel programming models

and run-time resource management techniques to exploit the features of manycore processor architectures. To this purpose, a proper Consortium has been set up to gather the required expertise in the areas of system/application software and computing architectures. The 2PARMA project focuses on the definition of a parallel programming model combining component-based and single-instruction multiple-thread approaches, instruction set virtualisation based on portable bytecode, run-time resource management policies and mechanisms as well as design space exploration methodologies for Manycore Computing Fabrics.

### COMPLEX project

The primary objective of COMPLEX is to develop an innovative, highly efficient and productive design methodology and a holistic framework for iteratively exploring the design space of embedded hardware/software (HW/SW) systems.The COMPLEX framework approach is as follows:

Highly efficient and productive design methodology and a holistic framework for iteratively exploring embedded HW/SW applications.

Augmentation of well established ESL tools enabling performance & power aware virtual prototyping from a combined HW/SW perspective.

Multi-objective co-exploration in combination with fast simulation and assessment at the earliest instant in the design cycle.

### PHARAON project

The PHARAON project targets the development of two different sets of techniques and tools, both aiming at best exploiting the low-power capabilities of modern multi-core processors, tackling both the programming and runtime power management challenges mentioned above. The first set will directly impact the design of the application. The objective is to assist the designer in finding the most adequate software architecture taking into account hardware constraints. To do so, tools will be capable to evaluate the parallel structure of an application and propose improvements. A tool will also be capable to handle communications between different processors and generate the multi-processor embedded code. The second set of techniques and tools will impact the runtime behaviour of the application. The objective is to adapt the performance of the platform, (frequency & voltage for example) in order to consume only the required energy. A reconfiguration system and a low power scheduler will be integrated with other run-time components on top of the platform.

UNIBO ETZH EPFL STM

### Pro3D Project

Pro3D is an example of project whose scope almost completely overlaps with the Artist-design MPSoC Design activity and whose participation is almost complelely coming from Parnters and associate partners of Artist-design.

PRO3D will innovate in both hardware and software technologies and demonstrate the effectiveness of manycores by an integrated and concerted effort in key aspects of hardware and software design. The uniqueness of this proposal stems from the experience of the partners in various aspects of manycore design that need to be addressed concurrently.

The key outcome of PRO3D will be a holistic system design methodology to bring a drastic improvement of productivity to reduce cost development and time to market for future embedded computing. In particular,

PROD3D will: i) develop a system software flow that can operate transparently on parallel manycore platforms;

ii) develop formal methods for software design guaranteeing the composability and correct operation of both

Hardware and software; iii) explore the impact of 3D integration for new computing architectures iv) Extend the

software-flow to 3D-stacked manycores.

*-- Changes wrt Y3 deliverable --*

*Some EU FP7 projects has been added to contribute to the NoE (e.g. CERTAINTY)*

# 4. Internal Reviewers for this Deliverable

- **Raphael David (CEA**)
- **Petru Eles (LIU)**