IST-004527 ARTIST2
Network of Excellence
on Embedded Systems Design

Activity Progress Report for Year 4

JPIA-Platform
System Modeling Infrastructure

Clusters:

Execution Platforms

Activity Leader:

Prof. Jan Madsen (Technical university of Denmark)
www.imm.dtu.dk/~jan

Policy Objective (abstract)

Integrate ongoing research efforts on infrastructure modelling. This would replace prototyping hardware to reduce the cost and time required for designing embedded systems. This activity is strategic for providing one angle in tackling the growing complexity of embedded systems.

Table of Contents

1 Overview of the Activity

1.1 ARTIST Participants and Roles

Prof. Petru Eles – ESLAB, Linköping University (Sweden)
Areas of his team's expertise: models for communication intensive distributed systems, power models, analytic performance estimation.

Prof. Dr. Rolf Ernst – IDA, TU Braunschweig (Germany)
Areas of his team's expertise: model infrastructure for performance analysis of heterogeneous systems.

Prof. Luca Benini – Micrel Lab, University of Bologna (Italy)
Areas of his team's expertise: providing models for estimation of non-functional properties.

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.

1.2 Affiliated Participants and Roles

Dr. Roberto Zafalon – STM (Italy)
Areas of his team's expertise: requirements and use of platform.
Salvatore Carta - University of Cagliari (Spain)
Working on MPSoC Middleware (task migration support) using our infrastrcucture
Dr.Magnus Hellring – Volvo (Sweden)
Areas of his team's expertise: requirements analysis.
Bjørn Sand Jensen – Bang & Olufsen ICEpower (Denmark)
Areas of his team's expertise: chip design for audio signal procesing
CTO Rune Domsteen – Prevas (Denmark)
Areas of his team's expertise: platform design for embedded systems

1.3 Starting Date, and Expected Ending Date

Starting date: September 1st, 2004

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

1.4 Baseline

A key research and research integration enabler is a scalable and realistic modelling platform which is abstract enough to provide complete system representations and some form of functional models even for billion-transistor future systems, while at the same time providing the needed flexibility for modelling a number of different embodiments (e.g. multi-processors, homogeneous and heterogeneous, reconfigurable, etc.).
The main objectives are:

- Integrate ongoing research efforts on infrastructure modelling
- Replacing prototyping of hardware

- Reducing the cost and time required for designing embedded systems
- Tackling the growing *complexity* of embedded systems

1.5 Problem Tackled in Year 4

The aim of the System Modeling Infrastructure activity is to create a system-level model in which designers can model and analyse complex heterogeneous embedded systems, and in particular explore:

- Consequences of different mappings of tasks to processors – software or hardware
- Effects of different RTOS selections – scheduling, synchronization and resource allocation policies
- Effects of different Network topologies and communication protocols.
- Effects of different component selections – processor, memory and memory hierarchy, IO buffers, co-processors

As this is a very complex problem, we are addressing it through two different approaches:

- Simulation-based modeling, in which we want to integrate cycle-true models with transaction level and abstract models in order to be able to do cross-layer and –level modeling and analysis.
- Formal modeling, in which we want to integrate different formalisms in order to be able to perform worst-case and response time analysis as well as schedulability analysis.

The major focus of Year4 has been on:

- Continue the investigation of design paradigms for MPSoC architectures. In particular by applying analysis and optimization methods used for distributed systems to tackle problems of MPSoC architectures.
- Continue the simulation-based and analytical methods for analysis of distributed real-time systems. In particular to use the simulation-based approach to judge the pessimism of analytical approaches.
- Extending the modelling and analysis of fault tolerant distributed systems. One aspect is to incorporate soft real-time tasks into the analysis methods for hard real-time tasks. Another aspect is to capture the reliability of the distributed system.
- Extending the development and integration of formal-based models. Coupling modular performance analysis and timed automata, and extending the MPSoC model based on timed automata to cope with systems of realistic sizes of both execution platform and application.
- Exploite and extend the simulation-based modeling for dynamically reconfigurable architectures, in order to explore different run-time management strategies.
- The usage of the models and extensions based on feedback from using the models. Most of the model usage is reported in the other activity reports from the Execution Platforms cluster.

1.6 Comments From Year 3 Review

1.6.1 Reviewers' Comments

The comments on the Y3 deliverable are positive: *ACCEPTED*
The document is well written. It clearly describes the achievements in year 3, the dissemination

activities and the integration activities. The list of publications indicates significant technical progress and interaction among the cluster members. The future work plans are appropriate and consistent with the roadmap.

NB: Paragraph 2 of section 5.4.5 of the Reviewer's Report (version 2.0 sent May 28th 2008) seems to be out of place (perhaps a cut and paste error). It asks about HRT, which is not related to our topic: "System modelling infrastructure".

1.6.2 How These Have Been Addressed

No particular comments related to this activity have been given. However, we have continued and intensified the activities on fault tolerance, which was introduced in year 3 based on the general comments given by the reviewers.

2 Summary of Activity Progress

2.1 Previous Work in Year 1

2.1.1 Simulation-based modeling (month 0 – 6)

MPARM (UoB)

UoB has devoted quite a lot of effort to developing a simulation environment for multi-processor systems on chip, called MPARM, and augmenting its modelling capabilities so to be able to simulate real-life systems. The work initially addressed the integration of IP cores into the simulation environment, and therefore required to cope with issues such as IP core wrapping and standard interfacing, synchronization with the system simulation engine, development of non-functional models (e.g., power consumption) and integration of such models in the functional simulation. Moreover, inter-processor communication and synchronization mechanisms have been developed, modelled and explored via functional simulation. The work on IP core integration, communication and synchronization has led to a multi-processor system-on-chip simulation environment with unprecedented modelling capabilities, which can be effectively deployed for design space exploration at a high level of accuracy. In fact, the cycle-accuracy of all component models has been preserved throughout the development and integration process.

MPARM existed prior to ARTIST2, but it has been extended and refined during the first year of ARTIST2. The work contributed by ARTIST2 is the development of the entire traffic tracing and the transaction extraction facilities, which have been instrumental to integrate the traffic generators (see next paragraph) with MPARM. On the modelling side, ARTIST2 has contributed with the development of a few communicaiton intensive benchmarks.for comparing different NoC technologies and strategies.

Traffic generators (UoB, DTU)

Traditionally, synthetic traffic generators have been used to overcome the more realistic development scenarios in the industry, where the parallel development of components may cause IP core models to be still unavailable when tuning the communication architecture. However, target applications increasingly present non-trivial execution flows and synchronization patterns, especially in presence of underlying operating systems and when exploiting interrupt facilities. This property makes it very difficult to generate realistic test traffic. For this reason, UoB and DTU have established a cooperation to realistically render SoC traffic patterns with interrupt awareness. The proposed methodology was extensively validated by showing cycle-accurate reproduction of previously traced application flows.

The entire work on developing, integrating and testing the traffic generators are contributed by ARTIST2. This has been done in a close cooperation between UoB and DTU.

ARTS (DTU)

ARTS is a system-level heterogeneous multiprocessor System-on-Chip (MPSoC) modeling framework which allows for designer-driven design-space exploration of heterogeneous MPSoC platform architectures and co-exploration of cross-layer dependencies. In particular, the consequences of different mappings of tasks to processors – software or hardware, the effects of different RTOS selections – scheduling, synchronization and resource allocation

policies, and the effects of different network topologies and communication protocols. As both ARTS and MPARM are based on SystemC and interface components using OCP, initial tempts to interface the two models have been taken.
ARTS existed prior to ARTIST2, but it has been extended and refined during the first year of ARTIST2. The work contributed by ARTIS2 are the development of an OCP-based interface between computaion and communication components. This has required a major rewrite of the modelling core of ARTS as well as an extension of the basic ARTS model to include IO tasks, modelling device drivers, and IO devices, modelling the interface devices.

2.1.2 Formal modeling (month 0 – 6)

Response time analysis (TU Braunschweig)

The aim has been to consider remote memory accesses in worst-case response times. In our extended task model, multiple remote transactions during the task's execution are taken into account. This extension allows the straightforward modelling of many applications and increases the applicability of formal analysis methods to many real-world architectures containing shared memory.

The task model extension that allows to include the effect of remote memory accesses on worst-case response time analysis is entirely contributed by ARTIST2.

Sensitivity analysis (TU Braunschweig)

Sensitivity analysis and flexibility optimization is used to reduce the design risk of critical components and increase design robustness. Sensitivity analysis allows the system designer to evaluate the flexibility of a given system, and thus to quickly assess the system-level impact of changes in performance properties of individual hardware and software components. If for example, the integration of a supplied IP component into the system in its original configuration results in a non-working system, the designer can easily determine if the reconfiguration of the system is possible so that all system constraints are satisfied.

The Artist2 contribution to sensitivity analysis is the formal framework to define sensitivity and robustness measures which can then be used in design space exploration and robustness optimization.

2.1.3 Simulation-based modeling (month 6 - 12)

MPARM (UoB)

The system modelling effort in MPARM has not only been concerned with the hardware architecture, but also the software infrastructure. A middleware layer has been designed with the objective of abstracting software developers from low level implementation details such as memory map, address of memory mapped slave devices, management of synchronization and inter-processor communication mechanisms, shared memory allocation and de-allocation, etc. This work has been performed with the cooperation of associated member IMEC.

The work on the middleware layer performed in cooperation with IMEC is entirely contributed by ARTIST2.

Trafiffic genrators (UoB, DTU)

The cooperation between UoB and DTU to establish a reactive traffic generator model has been continued in this period. The generator model has been extended to deal with more complex and more realistic events, such as OS-driven interrupt handling mechanism, and therefore mimic non-trivial execution flows and synchronization patterns.

The work on extending the traffic generators to handle interrupts and synchronization is entirely contributed by ARTIST2.

Distributed embedded systems for automotive applications (Linköping University, DTU)

The work of Linkoeping Univeristy aims at implementing a simulator for distributed embedded systems for automotive applications. The starting point for this work is the multiprocessor simulation environment (ARTS) developed at DTU. The work is performed in cooperation between the DTU and Linkoping groups. A student from DTU has made a short visit to Linkoping in June 2005.

The entire work on developing a simulator for distributed embedded systems for automotive applications based on ARTS is contributed by ARTIST2.

ARTS (DTU)

The work on the ARTS modeling framework has been extended to include more details of the platform. This has been done in order to target two different types of platforms: MPSoC particularly for multimedia applications, and wireless sensor networks. On the application side, DTU and Linköping University have started a cooperation aimed at extending ARTS to be able to simulate distributed embedded systems for automotive applications.

The cooperation between DTU and Linkoping is done entirely within ARTIST2. The extensions related to multimedia applications, including the setup and experimentation wich a wireless multimedia terminal is contributed ARTIST2. The work done on extending the ARTS modeling framework to support modelling of wireless sensor networks has been done within the Hogthrob project, which is one of the main sources of fundings for the system modelling infrastructure activity.

2.1.4 Formal modeling (month 6 – 12)

Modelling power consumption with SymTA/S (TU Braunschweig)

Recently the SymTA/S tool was extended in cooperation with Prof. Sharon Hu, University of Notre Dame, USA, to model and analyse the power consumption of complex heterogeneous embedded systems. Power aspects represent, besides performance issues, a critical problem during implementation and integration of complex systems. Currently we are working on power optimization techniques (heuristic and stochastic) based on the developed power models.

SymTA/S existed prior to ARTIST2. The work on extending the basic model of SymTA/S to handle power consumption which has been done in a cooperation between TU Braunschweig and University of Notre Dame is contributed by ARTIST2.

Integration (TU Braunschweig, ETHZ)

Initial attempts to integrate the event-stream formalism from TU Braunschweig and the real-time calculus from ETHZ have been established. The real-time calculus has been embedded into the SymTA/S front-end, allowing designers to expres both models within the same environment.

The cooperation between TU Braunschweig and ETHZ on embedding the real-time calculus within the SymTA/S front-end is entirely contributed by ARTIST2.

2.2 Previous Work in Year 2

Simulation platform for distributed embedded systems (University Linköping)

A simulation environment is designed and implemented for distributed real-time systems such as those used in automotive applications. The ARTS environment, developed at DTU and targeting System on chip applications, has been used as a starting point by the Linköping team.

In Year2 the following work has been done:

- The implementation of the environment has been finalized and new protocols, such as Flexrey have been implemented;
- Theoretical investigation regarding anomalies and sensitivity in distributed real-time systems has been performed, with results that will help to improve the efficiency of the simulator in detecting close to worst case behavior. This is important when using the simulator for evaluation of the pessimism of certain schedulability analysis approaches. This work is done in interaction with the Braunschweig group.
- Implementation of real-life applications from our industrial partners at Volvo.

First publication is planned for the next year.

Modeling and response-time/buffer analysis for NoC (University Linköping)

The Linköping group has developed a system model, based on which worst case response times and worst case buffer need for hard real-time applications implemented on NoCs can be calculated. On top of this analysis approach, an optimization tool for buffer space minimization has been implemented, for real-time NoC applications.

Modelling and formal timing analysis of shared memory accesses (TU Braunschweig)

We have continued to investigate design paradigms of MPSoC architectures. As opposed to distributed systems, a common feature here is the use of a shared memory that is accessed from each processor, introducing conflicts on the memory and interconnects. System designers often implement latency-hiding techniques to reduce the effect of waiting for data, by allowing frequent context switches to tasks that are ready.

We have systematically identified dependencies in such systems that have an influence on design properties such as end-to-end delays. Using this in [SIE06], we were able to show that the technique for latency hiding can bear unwanted results for critical worst-case response time scenarios.

We have further investigated formally the timing of multiple coinciding memory accesses. Previous approaches had to assume a worst-case timing for each individual memory access. Due to large timing variations, this leads to a large deviation of analysis result and actual

behaviour. In [SISE06], we presented a new way to express and calculate total latency of multiple events with much higher accuracy, leading to improve worst-case response time estimates.

Integration of formal SDF analysis techniques into the SymTA/S framework (TU Braunschweig)

Standard event models represent key integration aspects and hide complexity of local scheduling analysis algorithms. Thus, they are a suitable abstraction to integrate different models of computation into the SymTA/S framework. Recent work at IDA has produced a methodology to embed the analysis of SDF Graphs [Lee/Messerschmitt] into the SymTA/S framework (paper submitted for review at DATE07).

Integrating SDF models into the SymTA/S framework required corner-case evaluation of SDF graphs to construct event models describing their timing behaviour. Also, notions for path related metrics like latencies were defined and algorithms for computing their upper and lower bounds were proposed.

SDF Graphs are especially suited for describing data transforming applications like filters. Integrating their analysis into the SymTA/S framework significantly enlarges its application domain and improves the analysis results i.e. in the field of filter applications.

Multi-dimensional sensitivity analysis (TU Braunschweig)

The robustness of an architecture to changes is a major concern in embedded system design. Robustness is important in early design stages to identify if and in how far a system can accommodate later changes or updates or whether it can be reused in a next generation product. Robustness can be expressed as a "performance reserve", the slack in performance before a system fails to meet timing requirements. This is measured as design sensitivity.

Due to complex component interactions, resource sharing and functional dependencies, one-dimensional sensitivity analysis [RJE05] cannot cover all effects that modifications of one system property may have on system performance. One reason is that the variation of one property can also affect the values of other system properties requiring new approaches to keep track of simultaneous parameter changes.

Therefore, TU Braunschweig developed a heuristic and a stochastic approach for multi-dimensional sensitivity analysis [RHE06]. The heuristic approach is a divide-and-conquer like algorithm, which uses parameter specific heuristics to prune the search space. It is applicable to two dimensional search spaces. The stochastic approach is based on evolutionary search spaces and uses tabu search to bound the region containing the sought-after sensitivity front separating working and non working system configurations. It is applicable to search spaces of arbitrary dimension.

MPARM interface with Lisatek (University of Bologna)

New processor models have been included. The most important extension in this area is the integration with the SystemC models generated by the Lisatek suite developed by AACHEN. Any processor modeled in LISA can now be integrated as add-on core in the MARM platform. A standardized transaction-level interface has been defined for core embedding.

MPARM memory models (University of Bologna)

Models for external memory controllers (DRAM-DDRAM). The main memory interface is often the true performance bottleneck for many MPSoC platforms. Therefore significant effort has been devoted to the development of an accurate DRAM controller module, capable of several advanced communication-optimizations. The model has been integrated within the MPARM platform. Associate partner STmicroelectronics has provided the functional specification for the controller.

Traffic generator model (University of Bologna, Technical Univeristy of Denmark)

Applications running on MPSoC architectures increasingly present non-trivial execution flows and synchronization patterns, especially in presence of underlying operating systems and when exploiting interrupt facilities. These properties make it very difficult to generate realistic test traffic. Technical University of Denmark and University of Bologna have jointly developed a reactive traffic generator device capable of correctly replicating complex software behaviours in the MPSoC design phase. The approach has been validated by showing cycle-accurate reproduction of a previously traced application flow. The traffic models have been integrated in both the ARTS environment from Technical University of Denmark and the MPARM environment from University of Bologna.

ARTS modelling framework (Technical University of Denmark)

ARTS is a SystemC-based abstract system-level modelling and simulation framework, which allows MPSoC designers to model and analyze the different *layers*, i.e., application software, middleware and platform architecture, and their interaction prior to implementation. In particular, ARTS provides a simulation engine that captures *cross-layer* properties, such as the impact of OS scheduling policies on memory and communication performance, or of communication topology and protocol on deadline misses. The ARTS framework was demonstrated at the University Booth at the DATE07 conference in Munich. As a result, ARTS has been made public available. The distribution consists of the framework and a tutorial. The results of this work was published at MASCOTS05 [MSM05] and an article has been submitted for the journal on Design Automation for Embedded Systems.

A web-link to the downloadable ARTS framework is http://www.imm.dtu.dk/arts .

Toolbox for Modular Performance Analysis method of ETHZ (ETH Zurich)

The analytic performance analysis model for distributed embedded systems and multiprocessing devices has been refined and discussed together with other partners. A major event has been the Distributed Embedded Systems workshop in Leiden and the Execution Platform Meeting in Bologna. As a result, we decided to implement the basic mathematical tools of Real-Time Calculus in form of a Matlab toolbox. The aim is to foster even more integration in the future as now other groups will be able to apply and incorporate analytic methods easily. The first version of the toolbox is available, including documentation and a tutorial.

It will be used to integrate Symta/S and the modular performance analysis method in the next year of ARTIST2. A web-link to the toolbox is http://www.mpa.ethz.ch/Rtctoolbox/Overview .

Combining simulation and formal analysis for performance analysis (ETH Zurich)

Collaboration with Francesco Poletti and Luca Benini at University of Bologna

In this activity, we developed a new, compositional performance evaluation method for embedded systems. The new method combines existing approaches for system-level performance analysis, namely MPA a formal method and MPSim a simulation-based approach. To enable this combination, we defined the interfaces needed between the different performance evaluation methods. As a core of the approach, we propose a method to generate simulation stimuli from analytical models. In addition, we introduced a measure to assess the quality of a generated simulation trace with respect to its analytical description. In order to show the applicability of this new approach for performance evaluation, we implemented an example system for such a combined performance evaluation consisting of a multiprocessor system-on-a-chip. It is based on existing models for simulation and analytical models extended by the needed interfaces for the combination, including an implementation of the simulation trace generation algorithm. This combined model was then used for a case study of an application running on a multiprocessor system.

To achieve the results described above, several physical and phone meetings were held to coordinate the joint effort and to discuss future directions of this activity. The following two publications [KBPT06] and [KT06] describe the results of the joint activity.

2.3 Previous Work in Year 3

NoC Emulation Framework (UNIBO, EPFL)

NoC-based MPSoCs involve new and critical design challenges, such as the design of network interfaces and protocols to provide reliable on-chip communication to transport the data of the cores. Also, the selection of suitable custom topologies of switches for the applications of the target MPSoC is critical to provide the needed low-latency at the physical interconnection layer to transport the data of the cores. All these challenges require a very time-consuming and error-prone design and tuning process of on-chip interconnects to design power-efficient and high-performance MPSoC.

UNIBO developed, in cooperation with EPFL a combined hardware-software NoC emulation framework, which shows how flexible NoC emulation can be used as a powerful design tool for tuning and functional validation of on-chip interconnections for MPSoCs. This emulation framework is implemented onto a *Field Programmable Gate Array* (FPGA) platform and has as one of its main novelties the utilization of the FPGA as an active element in the emulation control layer to speed up functional validation and to add flexibility to the NoC configuration exploration, instead of merely being the platform where the circuit is prototyped, as emulation is typically used.

The emulation framework is able to test actual physical realizations of NoCs on silicon up to four orders of magnitude faster than *Hardware Description Language* (HDL) simulators (see Figure 1), while preserving cycle accuracy. In addition, the flexibility of the emulation framework can be exploited to define a procedure to rapidly validate and tune NoC physical implementation characteristics (e.g., buffer size, topology of switches, size of inter-switches links, etc.) for real-life traffic patterns of software applications that can be executed in the target MPSoCs or various software scenarios (e.g., bursts lengths, average on chip communication load, etc).

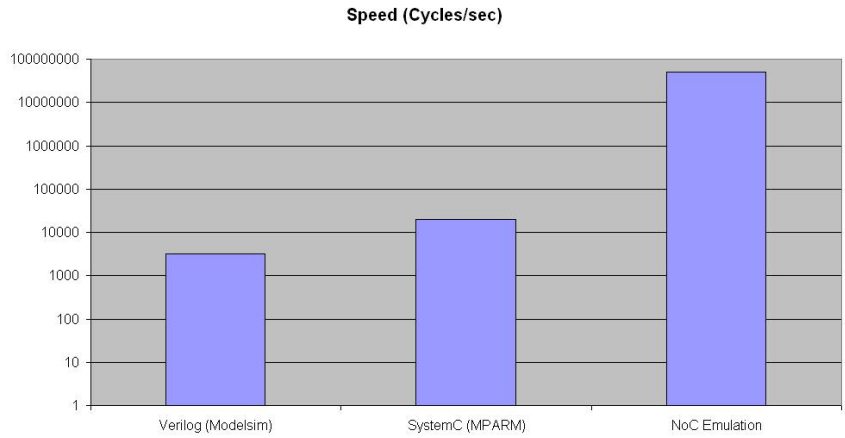

Figure 1: Speed comparison NoC-MPSoC simulation vs. emulation

Modeling and formal timing analysis of Multiprocessor Systems On Chip (TU Braunschweig)

TUBS have continued to investigate design paradigms of MPSoC architectures. As opposed to distributed systems, a common feature here is the use of a shared memory that is concurrently accessed from each processor, introducing conflicts on the memory and interconnects. System designers often implement latency-hiding techniques to reduce the effect of waiting for data, by allowing frequent context switches to tasks that are ready.

Building on previous work [SIE06] we have systematically investigated a realistic application with STMicroelectronics as an industrial partner.

The application was developed at the École Polytechnique de Montreal to run on the StepNP research platform. The involved round-robin scheduler could easily be integrated into our analysis engine. By conservatively considering the memory and bus congestion, this allowed to quickly model different architectural scenarios, and to predict corner case behavior which could not be identified in a simulation.

The work on the coupling a Synchronous Dataflow Graph based analysis with our compositional analysis approach has been presented at the DATE 2007 [SSE07]. It was received with great interest and has led to further cooperation with NXP Semiconductors, Eindhoven, NL.

Sensitivity Analysis and System Robustness Optimization for Complex Embedded Systems (TU Braunschweig)

TUBS have further extended their methods for sensitivity analysis and system robustness optimization.

As a result of HW/SW reuse, design data refinement or integration of components provided by different suppliers, the system designer must take into account that system properties, such as worst-case execution times, data rates, CPU clock rates, etc., are likely to be modified during the design process, or even later, during system life-cycle.

Our sensitivity analysis approaches can be used to compute, for the system properties subject to modification, the available slack with respect to an imposed set of constraints. Hence, any property modification carried out within the available slack interval guarantees that system feasibility is preserved. In many cases, the modification of a system property implies also the variation of other properties in the system. For such cases, we developed a multi-dimensional sensitivity analysis [RHE06].

In order to efficiently control performance and to ensure predictability, sensitivity analysis must be systematically integrated into the design flow of embedded systems. We, therefore,

proposed expressive robustness metrics for different assumptions and design scenarios, and showed how they can be efficiently considered throughout the whole design process. The proposed metrics are based on sensitivity analysis. At top level we distinguish robustness metrics for independent [HRE06] and dependent system properties [HRE07] w.r.t. system performance. For independent system properties the value of one system property does not have any influence on the admissible values for the other system properties. Contrarily, for dependent system properties the modification of one system property leads to more restrictions for the other system properties, i.e. their flexibility w.r.t. modifications decreases.
Performance characterization and system robustness become even more important if we assume that for complex application structures and dynamic scheduling policies, performance metrics, such as end-to-end latencies, response times, buffer sizes, etc., can easily exhibit unexpected non-monotonic behavior, a phenomenon known in literature as scheduling anomaly. In order to effectively cover such effects, we proposed a detailed scheduling anomaly analysis [RE06]. Our analysis can be used to find, on the one hand, system configurations with little design robustness, and on the other hand, to reveal additional performance reserves.

Fault-Tolerent Process Graph Model (DTU, LiU)

There is a lot of research in the area of system modeling and specification, and an impressive number of representations have been proposed. The system-level design tasks typically deal with sets of interacting processes. A process is a sequence of computations (corresponding to several building blocks in a programming language) which starts when all its inputs are available. When it finishes executing, the process produces its output values.

Researchers have used, for example, dataflow process networks (also called task graphs, or process graphs) to describe interacting processes, and have represented those using directed acyclic graphs, where a node is a process and the directed arcs are dependencies between processes. One drawback of dataflow process graphs is that they are not suitable to capture the different fault scenarios that can happen due to the occurrence of transient faults in a fault-tolerant application. For example, it can happen that the execution of some processes fails due to faults. By explicitly capturing such a failure in the process graph model, a more fine-tuned modeling and a tighter (less pessimistic) assignment of execution times to processes is possible, compared to traditional data-flow based approaches.

Together with Linkoeping University (LiU) DTU have proposed an extension to the process graph model, namely a “fault-tolerant process graph” model (FT-PG). In an FT-PG the fault occurrence information is represented as conditional edges, and thus the FT-PG captures all the fault scenarios that can happen during the execution of application [TVLSI]. We have shown how design transformations that introduce redundancy, such as re-execution and replication, can be applied on this model.

MOVES, Modeling and Verification of Embedded Systems (DTU, AAU)

One of the major challenges in designing an embedded system is to find a mapping of the application onto the execution platform, which effectively fulfills the non-functional requirements of the embedded system such as timing, memory usage, energy consumption, and other cost. A particular challenge is to model and analyse cross-layer dependencies, where the change of a property in one part of the system, e.g. scheduling policy, may impact the performance of another part of the system, e.g. deadline miss on another processor, and hence, the overall system performance. The ARTS simulation model developed by DTU during the first two years of ARTIST2, has been modeled using the semantics of timed automata and implemented in UPPAAL from AAU.

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

Modeling and Verification of Hardware Components (DTU)

As the complexity of chips grows, the methodology to build chips has to evolve. Today, chips are largely synthesized from high-level architectural descriptions that hide low-level details.
The majority of hardware designs are done using the most common hardware description languages, suvh as VHDL or Verilog. Both languages support high-level architectural descriptions, but allow hardware designers to incorporate low-level details in order to optimize for a particular hardware technology and directly synthesize using a restricted subset of the languages. However, chips may also be synthesized from software based models in much the same way as compilers produce executable code. Examples of such languages are Esterel, Lustre and Signal.
DTU have developed a language for hardware models based on the Gezel hardware description language developed and maintained by Virginia Tech, USA. The language depends on reasonably few, simple and clean concepts, and it strikes a balance between software and hardware concerns that suits the needs for a modern top-down approach to hardware design.

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

Simulation Platform for Dynamical Reconfigurable Systems (DTU)

One of the biggest challenges in reconfigurable system design is to improve the rate of reconfiguration at run-time by reducing the reconfiguration overhead. Such overhead comes from multiple sources, and without proper management, the flexibility of the reconfiguration can not justify the overhead cost. DTU have developed a flexible framework, called COSMOS, to model and simulate coprocessor-coupled reconfigurable systems. The framework is an extension to the ARTS framework developed by DTU during the first two years of ARTIST2. DTU have develpoed a novel real-time task model that captures the characteristics of dynamically reconfigurable systems' task in terms of initialization, reconfiguration and reallocation. DTU also propose a general model of coprocessor-coupled reconfigurable systems. The task and architecture models have been extended to facilitate the study of run-time resource management strategies. Based on this model, DTU have demonstrated how a simple "worst case" run-time system can be modelled in the COSMOS framework as a firmware to manage the application execution.

The COSMOS framework have been used to experiment with various combinations between the application and the architecture to gain a better understanding of the critical issues in reconfigurable architecture design. A set of experiments based on a MP3 task graph have been conducted.

MT-ADRES: Multithreading on Coarse-Grained Reconfigurable Architecture (DTU, IMEC)

To investigate the performance bottleneck and the scalability of the state-of-the-art datapath-coupled reconfigurable architectures, DTU and IMEC have studied the coarse-grained reconfigurable architecture ADRES (Architecture for Dynamically Reconfigurable Embedded Systems) developed by IMEC, Belgium. In order to improve task-level parallelism, they have proposed a method for multithreading on the ADRES architecture.

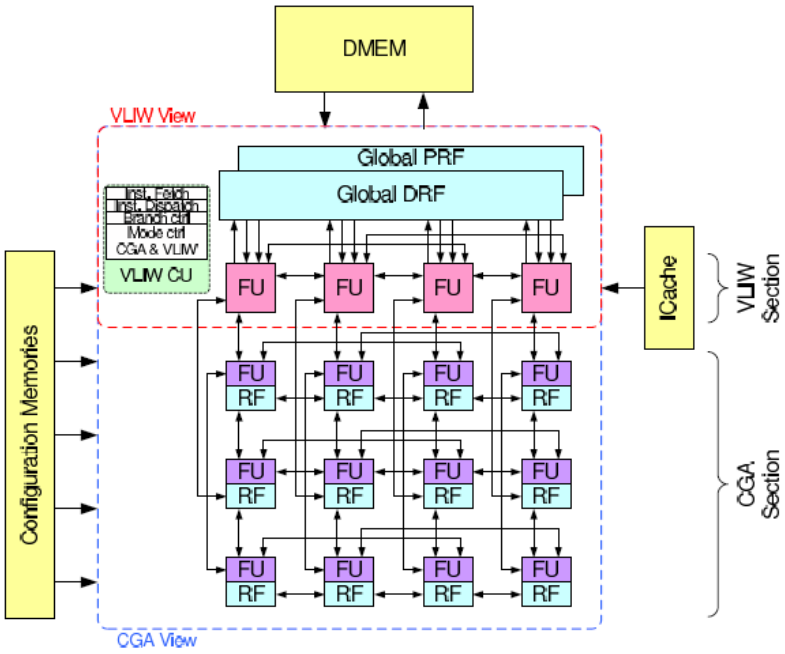

DTU and IMEC have proposed how the ADRES architecture can be extended to support multi-threading, and how the ADRES compilation tool flow needed to be extended to cope with multithreading. They have made an experiment running a dual-threaded MPEG2 decoder on a customized ADRES architecture to demonstrate that multi-threading is feasible for ADRES. Through the MPEG2 experiment they have discovered some design pitfalls that hinder the performance of the threaded ADRES, and discussed what technologies can further improve the performance of the multi-threaded ADRES.

MPA Model Integration (ETHZ, TUBS)

ETHZ has been combining its tool set (MPA – Modular Performance Analysis) with the system modelling infrastructure of other partners. Especially, there has been a deep integration between Symta/S and MPA. This not only entails converters between the two modeling formalisms but also investigations, when to use which formalism. In particular, it turns out that there are components of a design that can be much more accurately modeled by one or the other model. Besides the tool integration, there has been a paper published at CODES/ISSS that describes the obtained results.

In addition, there was an integration of the PISA multi-objective optimization framework into Symta/S as well as in tool to determine the robustness of a design at TU Braunschweig. Therefore, the exchange of tools and the integration between different modeling formalisms and too domains has been succesfully demonstrated. These goals heve been achieved by means of mutual visits, e.g. Simon Kuenzli (ETHZ) spend time at TU Braunschweig in October 2006.

Simulation platform for distributed embedded systems (LiU)

LiU have finalized their simulation platform for distributed embedded systems. Once the platform was available the efforts were concentrated into the following two directions:

1. Elaboration of a simulation methodology which allows to efficiently estimate the worst case response time of distributed real-time applications. In order to achieve an efficient simulation, two problems had to be solved:
   a) how to reduce the space of execution times to be explored;
   b) how to generate the next exploration point at a given moment of the simulation process? In other words, what exploration strategy to use.

   Simulations, if well conducted, can lead to tight lower bounds on worst-case response times, which can be an essential input at design time. Moreover, such a simulation methodology is very important in situations when the running application or the underlying platform is such that no formal timing analysis is available.
2. LiU have used the elaborated simulation platform two validate formal analysis approaches, by estimating their degree of pessimism. They have performed such an estimation of pessimism on two response-time analysis approaches for distributed embedded systems based on two of the most important automotive communication protocols: CAN and FlexRay.

Modeling and analysis for NoC communication (LiU)

The Linköping group has continued its work on the modeling and analysis of NoC platforms. In particular fault tolerance in the context of NoCs and transient faults has been addressed. Based on the elaborated modeling and analysis approach a system optimization methodology has been developed.

2.4 Final Results

2.4.1 Technical Achievements

The aim is to provide a scalable and realistic modelling platform which is abstract enough to provide complete system representations and some form of functional models even for billion-transistor future systems, while at the same time providing the needed flexibility for modelling a number of different embodiments (e.g. multi-processors, homogeneous and heterogeneous, reconfigurable, etc.).

The focus for year 4 was mostly to continue and extend the formal-based models, although some work on simulation-based models would be continued. For the formal-based models, the focus was to continue and extend the semantic model of SymTA/S to efficiently cover MPSoC architectures and to further investigate challenging timing issues in multiprocessor systems.

The work on models for the analysis and optimization of fault-tolerant embedded systems should be continued and further extended to capture fine-grained combinations of several fault-tolerance techniques. Two major efforts were planed for the models based on timied automata; 1) to continue the work on formalizing the ARTS model using timed automata based on UPPAAL, and in particular, to refine the formal model to address modeling and verification issues closer to the hardware layer of the execution platform. 2) to combine Modular Performance Analysis with timed automata based evaluation methods.
For the simulation-based models, the focus was to continue the work on extending the simulation-based model towards handling dynamically reconfigurable architectures, and iIn particular, to study different run-time resource management strategies. And for the simulation environment for distributed embedded systems, to do experimental evaluations to address the level of pessimism obtained through the formal methods.

The objectives have all been achieved. In the following, details are given for each sub-activity, listing a title and the partners involved in the sub-activity.

Modelling and Formal Performance Analysis of Multiprocessor Systems On Chip (TU Braunschweig)

TU Braunschweig has continued to investigate design paradigms of MPSoC architectures. As opposed to distributed systems, a common feature here is the use of a shared memory that is concurrently accessed from each processor, introducing conflicts on the memory and interconnects.

We have additionally investigated the options of applying the optimization techniques known from distributed multiprocessor systems with the analysis methodologies for multiprocessor systems on chip. The approaches turned out to be flexible enough to be easily adapted to exploring parameters of memory controllers and traffic shapers. With this, an optimal configuration for a realistic media processing application could be found, in which a careful balance between worst case latency and guaranteed throughput is required from the memory subsystem. The results of this experiment have been presented in [SHR+08].

With the investigation of a realistic application by the École Polytechnique de Montréal and STMicroelectronics as the supplier of the hardware models, we could also benchmark our analysis approach. Our conservative models have turned out to overestimate the worst simulated values by no more than 25% in the given system, which is a fair value given the increase in design confidence. The results will be published in [SNN+08].

Simulation-based and analytical methods for validation of distributed real-time systems (Linkoping)

We have finalized and improved our simulation-based and analytical platforms for validation of distributed embedded systems. One of the main directions was the evaluation of the simulation-based platform with regard to its capacity to be used for validation of time critical systems. Based on experimental results with the simulation-based platforms we were also able to draw interesting conclusions regarding the potential pessimism of formal approaches to validation of distributed real-time systems. The experiments were performed considering FlexRay and CAN-based distributed systems, which are of great interest for the automotive industry. Results were published at DATE 2008.

Integration of thermal models in a framework for energy efficient design of time constrained embedded systems (Linkoping)

This work is built on top of our research regarding energy efficient design of real-time systems. We have earlier developed several approaches to dynamic voltage selection (DVS), a technique which exploits the available slack times by reducing the voltage and frequency at which the processors operate and, thus, achieves energy efficiency.

It is known that high power densities achieved in current SoCs do not only result in huge energy consumption but also lead to increased chip temperatures. Growing temperature leads, among others, to an increase in leakage power and, consequently, energy, which, again, produces higher temperatures. Nevertheless, the temperature issue has been completely ignored in the DVS techniques for real-time embedded systems proposed in literature.

We have used state of the art temperature models for multicore chips and integrated them into a temperature aware dynamic voltage selection approach. We have shown that important energy savings can be achieved. Results have been published at DATE 2008.

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

LiU and DTU have continued their cooperation in the area of fault-tolerant real-time systems. The main focus was extending the system model in order to capture not only hard real-time applications but also soft real-time systems and, more generally, systems consisting of both hard and soft real-time tasks. The goal is to guarantee deadlines for the hard processes even in the case of faults, while maximizing the overall utility. We use time/utility functions to capture the utility of soft processes. Results have been published at DATE 2008.

DTU and LiU have extended their fault-tolerant process graph model for embedded systems to consider a combination of hardware and software fault-tolerant techniques. The model can capture re-execution of processes and hardened, i.e., more reliable, hardware processing elements. A System Reliability Analysis (SRA) technique has been proposed that calculates the reliability of the system considering: (a) the different hardening levels in hardware, (b) the re-execution levels in software (how many redundancies are there for a process), (c) the mapping, which decides what hardening-level will a process start from and (d) the scheduling, which decides how are the recovery slacks shared. Proposing such a SRA has been challenging because, due to the slack sharing, there are many complex combinations in which an application can fail.

MOVES, Modeling and Verification of Embedded Systems (DTU, AAU)

DTU has continued their work on formalizing the ARTS simulation model and to make it usable for designers early in the design process. In order to support designers of industrial applications, the timed-automata model is hidden for the user, allowing the designer to work directly with the abstract system-level model of embedded systems. The designer provides an application consisting of a set of task graphs, an execution platform consisting of processing elements interconnected by a network and a mapping of tasks to processing elements. The system model is then translated into a timed-automata model which enables schedulability analysis as well as being able to verify that memory usage and power consumption are within certain limits. In the case where a system is not schedulable, the tool provides useful information about what caused the missed deadline. DTU does not propose any particular methodology for design space exploration, but provide an analysis framework, MoVES, where embedded systems can be modelled and verified in the early stages in the design process.

Thus, the MoVES Analysis Framework provides tool support for system designers to explore alternatives in an easy and efficient manner.
An important aspect in the design of MoVES is to provide an experimental framework, supporting easy adaptability of the ”core-model” to capture energy and memory considerations, for example, or to experiment with, say new principles for task scheduling and allocation. Furthermore, the MoVES Analysis Framework is equipped with different underlying UPPAAL models (some of which have been developed together with AAU), aiming at efficient verification in various situations. For the moment DTU is operating with the following underlying models for

- schedulablity analysis in connection with worst-cases execution times only,
- schedulablity analysis for the full core model (including best- and worst-case executing times),
- schedulability analysis addressing memory and energy issues as well, and
- schedulability analysis for the full core model on the basis of stop-watch automata. This analysis approach is based on over approximations, but it has provided exact results in the experiments carried out so far and it appears to be the most efficient Uppaal implementation.

Formal verification of design properties of hardware architectures (DTU)

DTU has continued its work on a formal language for hardware models based on the Gezel hardware description language developed and maintained by Virginia Tech, USA. The language depends on reasonably few, simple and clean concepts, and it strikes a balance between software and hardware concerns that suits the needs for a modern top-down approach to hardware design. The semantical domain of the language has been related to timed-automata using the UPPAAL system. They have demonstrated a formal verification of design properties of a few simple example circuits including the Simplified Data Encryption Standard (SDES) Algorithm and different algorithmic implementations of the Greatest Common Divisor. Verification guarantees properties of the underlying algorithm, e.g. correct output for any given input, as well as other properties such as upper limits on the number of clock cycles for the algorithm to stabilize with a given input and upper limits on the number of register updates, to serve as an indicator of energy consumption.

Simulation Platform for Dynamical Reconfigurable Systems (DTU)

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. Through a number of design space exploration experiments, they have pinpointed the critical design issues in the reconfigurable architecture study and analyze their impact on the architecture performance.

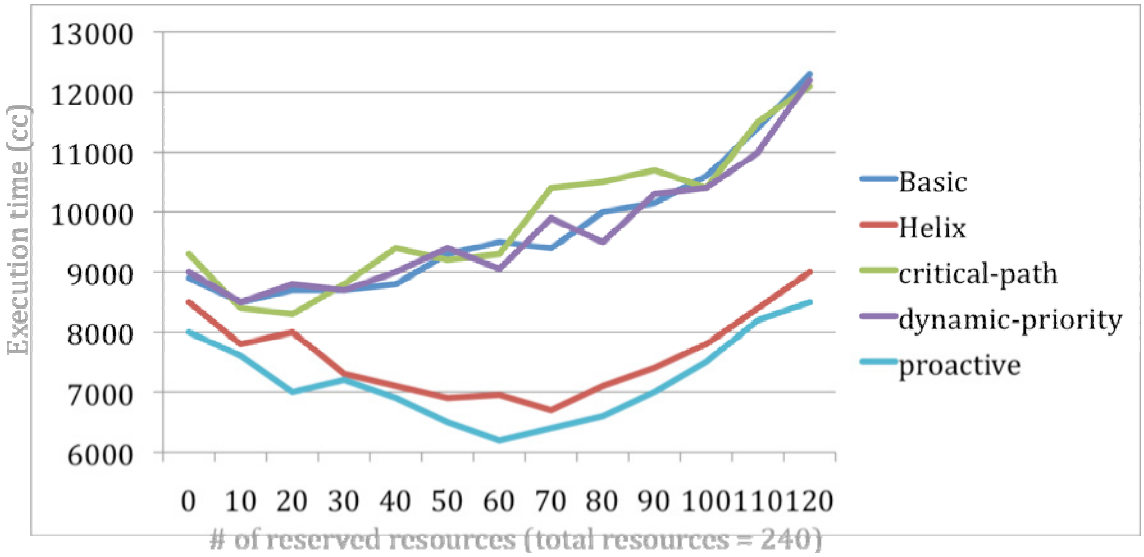

Analysis tools for embedded systems (DTU, Oldenburg)

This activity addresses the problem of establishing decidability and model-checking results for fragments of interval logics. The aim is, in particular, to establish efficient tools for the analysis of real-time properties of embedded systems, and, in general, tool support for the analysis of more general resource constraints of embedded systems. The activity is a co-operation with Prof. Martin Fränzle, Oldenburg University.

Composing analysis frameworks

As proposed, ETH Zurich coupled two different analysis frameworks, i.e. MPA (Modular Performance Analysis) and Timed Automata. This work has been done with the affiliated partner NUS (P.S. Thiagarajan). To this end, one of his PhD students visited ETH for a duration of 6 months. The results have been published and are now the basis for further work. In addition, ETH Zurich has been leading an activity where several partners from ARTIST2 have been involved which compared various abstraction mechanisms used in the analysis of distributed embedded systems. This work resulted in a joined conference and in a joined journal publication.

Modeling patterns for performance analysis (TU/e, ETHZ)

Modeling patterns for specification of applications and platforms and the automatic transformation to executable models for performance analysis. The key idea is to decouple the application and platform specification from the specific formal modeling and analysis tools. (TU/e work)

Comparison of different Y-chart based approaches (Metropolis, MPA and SHE/POOSL) for design-space exploration (combined work of University of Montral (Canada), TU/e and ETH Zurich)

2.4.2 Individual Publications Resulting from these Achievements

TU Braunsweigh

[RHE08] Razvan Racu and Arne Hamann and Rolf Ernst. "Sensitivity Analysis of Complex Embedded Real-Time Systems." In *Real-Time Systems*, Volume 39, pp 31-72, 2008.

[SHR+08] Simon Schliecker and Arne Hamann and Razvan Racu and Rolf Ernst. "Formal Methods for System Level Performance Analysis and Optimization." In *Proc. of the Design Verification Conference (DVCon)*, San José, CA, February 2008.

Linköping University

[SRE08] S. Samii, S. Rafiliu, P. Eles, Z. Peng, "A Simulation Methodology for Worst-Case Response Time Estimation of Distributed Real-Time Systems," Design, Automation, and Test in Europe (DATE 2008), Munich, Germany, March 2008.

[BAE08] M. Bao, A. Andrei, P. Eles, Z. Peng, "Temperature-Aware Voltage Selection for Energy Optimization," Design, Automation, and Test in Europe (DATE 2008), Munich, Germany, March 2008.

Technical University of Denmark

[MHK08] Jan Madsen, Michael R. Hansen, Kristian S. Knudsen, Jens E. Nielsen , Aske W. Brekling , System-level Verification of Multi-Core Embedded Systems using Timed-Automata, to appear in the proceedings of the 17th IFAC World Congress, Seoul, Korea, July 2008.

[BHM08]A. Brekling, M.R. Hansen, and J. Madsen. A timed-automaton model for multi-processor system-on-chips. Journal of Logic and Algebraic Programming, 2008.

[BHM08b] A. Brekling, M.R. Hansen, and J. Madsen. Formal Verification of Design Properties of Hardware Architectures. Poster at DATE 2008, UnivBooth

[WHM08] Kehuai Wu, Esben Rosenlund Hansen, Jan Madsen, Towards Understanding and Managing the Dynamic Behavior of Run-Time Reconfigurable Architectures, to appear in the proceedings of the Conference on Engineering of Reconfigurable Systems and Algorithms (ERSA'08), July 2008.

[BHM07] Brekling, A., Hansen, M. R., Madsen, J., Hardware Modelling Language and Verification of Design Properties, Nordic Workshop on Programming Theory, pp 49-51, 2007.

[HMB07] Hansen, M. R., Madsen, J., Brekling, A., Semantics and Verification of a Language for Modelling Hardware Architectures, Chapter in Formal Methods and Hybrid Real-Time Systems edited by Cliff B. Jones Zhiming Liu and Jim Woodcock, LNCS 4700, pp. 300-319, 2007

ETHZ

[CMR07] Samarjit Chakraborty, Tulika Mitra, Abhik Roychoudhury, Lothar Thiele, Unmesh D. Bordoloi, Cem Derdiyok: Cache-Aware Timing Analysis of Streaming Applications 19th Euromicro Conference on Real-Time Systems (ECRTS), Pisa, Italy, pages 159 - 168, July, 2007.

[HaT07] Wolfgang Haid, Lothar Thiele: Complex Task Activation Schemes in System Level Performance Analysis Proc. 5th Intl Conf. on Hardware/Software Codesign and System Synthesis (CODES+ISSS 2007), ACM Press, Salzburg, Austria, pages 173-178, October, 2007.

Tu/e

[FVT08] O. Florescu, J. Voeten, B. Theelen and H. Corporaal. Patterns for Automatic Generation of Soft Real-Time Models. Invited article, Simulation Journal, Special Issue on Multi-Paradigm Modeling: Concepts and Tools.

2.4.3 Interaction and Building Excellence between Partners

The following is a list of the major collaborations between partners in the System Modeling Infrastructure activity as well as with partners outside the activity which has contributed to creation of excellence between partners.

Activities within the activity:

- LiU-DTU: The co-operation on modelling, analysis and optimization of fault-tolerant distributed embedded systems has been continued and intensified in year 4. Several visits between DTU and LiU have taken place, and the co-operation has resulted in 6 joint publications.
- TUBS-ETHZ: Joint work on comparing performance analysis methods. The result is published in 2 joint publications.
- Tu/e-ETHZ: Joint work on approaches for system design based on Y-charts. The result is published in 1 joint publication.
- TUBS-LiU: Together with Rolf Ernst (TUBS), Petru Eles (Linköping) is co-chairing the dissertation of Razvan Racu (TUBS). During his dissertation, collaborations were conducted on several topics on schedulability analayis of multiprocessor systems.

Activities with partners outside the activity

- TUBS-Montreal-ST: Joint work on investigating a realistic application obtained from Ecole Polytechnique de Montreal and a model of the hardware platform from STMicroelectronics. Benchmark results are published in 1 joint publication.
- ETHZ-NUS: There was a close interaction between ETH Zurich and NUS Singapore during the work on coupling different analysis methods, including a long PhD visit. The comparison between the different evaluation methods involved the groups at TUBS, ETHZ and the Universidad de Cantabria. The result is published in 1 joint publication.
- LiU-Volvo: Joint work on formal method-aided simulation for model validation for embedded systems. The result is published in 1 joint publication.
- DTU-IMEC: Has finished their collaboration on extending the ADRES architecture and toolchain to support multitask programming. A joint journal publication has been published.
- DTU-Oldenburg: Prof. Martin Franzle is a guest professor at DTU. Martin Franzle has visited DTU several times for one or two weeks. He participated and lectured in the PhD course on Automated Formal Verification for Embedded Systems organized by DTU at DTU, June 4-12, 2007. Michael R. Hansen from DTU has vistited Oldenburg several times.
- DTU-AAU: Joint work on modeling the ARTS framework using the timed automata semantics of UPPAAL and extension to handle detailed hardware aspects. Several visits between DTU and AAU have taken place. Results have been published by DTU in 4 papers.

2.4.4 Joint Publications Resulting from these Achievements

1. [IPP+09] Viacheslav Izosimov, Ilia Polian, Paul Pop, Petru Eles, Zebo Peng, “Analysis and Optimization of Fault-Tolerant Embedded Systems with Hardened Processors”, Submitted to: DATE Conference, 2009.
2. [SNN+08] Simon Schliecker and Mircea Negrean and Gabriela Nicolescu and Pierre Paulin and Rolf Ernst. "Reliable Performance Analysis of a Multicore Multithreaded System-On-Chip." In *Proc. 6th International Conference on Hardware Software Codesign and System Synthesis (CODES-ISSS)*, Atlanta, GA, October 2008.
3. [KEP08] D. Karlsson, P. Eles, Z. Peng, "Model Validation for Embedded Systems Using Formal Method-Aided Simulation," IET Computers & Digital Techniques journal (accepted for publication).
4. [IPE08] V. Izosimov, P. Pop, P. Eles, Z. Peng, "Scheduling of Fault-Tolerant Embedded Systems with Soft and Hard Timing Constraints," Design, Automation, and Test in Europe (DATE 2008), Munich, March 2008.
5. [WMK08]Wu, K., Madsen, J., Kanstein, A., Mladen, B., *MT-ADRES: Multithreading on Coarse-Grained Reconfigurable Architecture*, to appear in the Intel. Journal of Electronics (IJE), Volume 95, Issue 7, July 2008. Page(s): 761-776.
6. [PPE08] T. Pop, P. Pop, P. Eles, Z. Peng, A. Andrei, "Timing Analysis of the FlexRay Communication Protocol," Real-Time Systems Journal, Volume 39, Numbers 1-3, August, 2008, pp 205-235.
7. [PPE08b] T. Pop, P. Pop, P. Eles, Z. Peng, "Analysis and Optimisation of Hierarchically Scheduled Multiprocessor Embedded Systems," Intl. Journal of Parallel Programming, Volume 36, Number 1, February, 2008, pp. 37-67.
8. [PIE08] Paul Pop, Viacheslav Izosimov, Petru Eles, and Zebo Peng, “Design Optimization of Time- and Cost-Constrained Fault-Tolerant Embedded Systems with Checkpointing and Replication”, In IEEE Transactions on Very Large Scale Integration (VLSI) Systems, 2008.
9. [EIP08] Petru Eles, Viacheslav Izosimov, Paul Pop, Zebo Peng, “Synthesis of Fault-Tolerant Embedded Systems”, In proceedings of the Design, Automation, and Test in Europe Conference, pp. 1117-1122, 2008.
10. [PWT08] Simon Perathoner, Ernesto Wandeler, Lothar Thiele, Arne Hamann, Simon Schliecker, Rafik Henia, Razvan Racu, Rolf Ernst, Michael González Harbour: Influence of Different Abstractions on the Performance Analysis of Distributed Hard Real-Time Systems Design Automation for Embedded Systems, Springer Science+Business Media, LLC, 2008.
11. [LTS08] Y-chart Based System Design: A Discussion on Approaches. J. Lapalme, B. Theelen, N. Stoimenov, J. Voeteb, L. Thiele and E Aboulhamid. Submitted to ACM Transactions on Embedded Computing Systems, special issue on model-driven embedded system design.
12. [PCT07] Linh Phan, Samarjit Chakraborty, P. S. Thiagarajan, Lothar Thiele: Composing Functional and State-based Performance Models for Analyzing Heterogeneous Real-Time Systems 28th IEEE Real-Time Systems Symposium (RTSS), Tucson, Arizona, US, December, 2007.
13. [PWT07] Simon Perathoner, Ernesto Wandeler, Lothar Thiele, Arne Hamann, Simon Schliecker, Rafik Henia, Razvan Racu, Rolf Ernst, Michael González Harbour: Influence of

Different System Abstractions on the Performance Analysis of Distributed Real-Time Systems ACM & IEEE International Conference on Embedded Software (EMSOFT), ACM Press, Salzburg, Austria, pages 193-202, October, 2007.

2.4.5 Keynotes, Workshops, Tutorials

Tutorial : Formal methods in system and MpSoC performance analysis and optimization (R. Ernst, S. Charkaborty, Hans Sarnowsky, Marco Bekooj, M Jersak)

DATE 2008
Munich, Germany – March 10, 2008

The tutorial provided an introduction to formal platform performance analysis covering the main communication and resource modelling techniques and their application to embedded systems and MpSoC. It included industrial applications and experiences, use cases from automotive design, which demonstrated how to acquire the necessary model data, an overview on predictable MpSoC platform sharing using service shaping concepts, as well as an introduction on how to combine state-based and functional models for MpSoC in a single analysis to improve modelling precision.

Mini-Keynote : Load level modelling (R. Ernst)
MpSoC Conference

St. Gerlach, The Netherlands, June 23-27, 2008.

Current ESL methods and tools for verification focus on run-time efficient simulation. Simulation requires executable code and is, therefore, not applicable to early design phases where the software code is not yet available or is still subject to changes. In such cases, load models that are known from schedulability and network analysis can be used. Load models are compatible to the activation rules of application models, such as event driven data flow graphs or time driven Simulink models. The main limitation of classical schedulability analysis is its focus on worst-case design. The presentation will outline an extension to the classical load model that captures task execution time and communication variations. That model can raise platform design to a next level of abstraction, thereby supporting a design process where application development is separated from software implementation. Using this model, analysis can also highlight sensitivity to software and application modifications.
http://www.mpsoc-forum.org/

Lecture: A Theory of Duration Calculus with Applications
PhD school on Domain Modeling and Duration Calculus
Shanghai, China - September, 2007

Michael R. Hansen from DTU has given a lecture with the title “A Theory of Duration Calculus with Applications” at the PhD school in Shanghai, China.

Summerschool: Advanced Digital Systems Design
Lausanne, Switzerland – 10-14 September, 2007

ETH Zurich has given part of a summer school/advanced course on ADVANCED DIGITAL SYSTEMS DESIGN. The participants are from industry and university. This way, results from the integrated view of embedded system design will be brought to a much larger community.

Workshop: Models of Computation and Communication
Eindhoven, The Netherlands – 3-4th July, 2008

As a follow-up of the Models of Computation and Communication (MoCC) at the ETH last year, the TU/e organized the MoCC2008 workhop. It brought together scientists from various areas, i.e. formal methods, hardware design and software architecture,
http://www.artist-embedded.org/artist/MoCC-2008.html

Workshop: CASTENESS
15.-18th of January 2008

ETH Zurich has been organizing and participating in the CASTENESS Workshop, see www.casteness.org. The workhop put together the expertise of various EU projects such as ARTIST2, SHAPES, AETHER. In addition, ETH Zurich has been given a tutorial on issues that have been investigated in the ARTIST2 context: Analytic Performance Estimation, Mapping Algorithms to Architectures, Scalable SW Construction. The workshop has been sponsored by ARTIST2.

Tutorial: EMSOFT
30th of September 2007

Lothar Thiele from ETHZ has been given a tutorial at EMSOFT, the major conference in the area of embedded software. It covered methods for performance analysis of distributed embedded systems and presented outcomes of the ARTIST2 project.

Invited talk:
17th International Federation of Automatic Control (IFAC) World Congress
Seoul, Korea – 6-11 July, 2008

Jan Madsen from DTU gave an invited talk on “System-level Verification of Multi-core Embedded Systems using Timed-Automata”. The talk presented the MoVES modelling framework, which allows for cross-layer modeling and verification, covering the application layer, middelware layer (RTOS), and hardware layer. The modelling framework allows the designer to verify the impact of execution platform and application mapping on the schedulability (meeting hard real-time requirements), power consumption and memory utilization, while taking communication into account. The modeling framework is implemented using timed-automata in UPPAAL.

Conference: Industrial Embbed Systems
3rd International Symposium on Industrial Embedded Systems (SIES)
Montpellier, France – 11-13 June, 2008

Anders Tranberg-Hansen from DTU gave a talk on “A Service Based Estimation Method for MPSoC Performance Modelling”. The talk presented an abstract service based estimation method for MPSoC performance modeling, which allows fast, cycle accurate design space exploration of complex architectures including multi processor configurations at a very early stage in the design phase. To illustrate the method, a small MPSoC system, developed at Bang & Olufsen ICEpower was modelled and performance estimates were produced for various configurations of the systemimplementation.

Demo: Formal Verification of Design Properties of Hardware Architectures
DATE 2008
Munich, Germany – March 10, 2008

DTU has given a demo of their tool for formal verification of design properties of hardware architectures at the DATE University Booth. In a 2 hour slot, the tool was presented and discussed with academic and industrial peopole participating in the DATE conference.

PhD-course: Automated Formal Methods for Embedded Systems
Lyngby, Denmark – 16-24 June, 2008

DTU has organized the second ARTIST2 sponsored PhD course on "Advanced Topics in Embedded Systems", that took place at IMM, DTU, Lyngby, Denmark, June 16-24, 2008. Lectures were given by ARTIST memebers from Oldenburg. The course was again a success and will be repeated in 2009.
http://www.artist-embedded.org/artist/Automated-Formal-Methods.html

Mini-keynote: Codesign
8th International Forum on Application-Specific Multi-Processor SoC (MPSoC)
St. Gerlach, The Netherlands, June 23-27, 2008.

Jan Madsen from DTU gave a talk on "*Adaptive Embedded Systems Challenges of Run-Time Resource Management".* The mini-keynote addressed some of the challenges of run-time resource management in adaptive embedded systems. The challenges come from the often very complicated interplay between the application, the application mapping, and the underlying hardware architecture. With MPSoC, we have the technology to design and fabricate dynamically reconfigurable hardware platforms. However, such platforms will pose new challenges to tools and methods to efficiently explore these platforms at run-time.
http://www.mpsoc-forum.org/

3 Milestones, and Future Evolution Beyond the NoE

3.1 Milestones

The milestones for year 4 defined at the end of year 3, have all been achieved as explained in the beginning of section 2.4.1. The following figure shows the evolution of the major activities of the activity on System Modelling Infrastructure. Boxes in red are simulation-based approaches, while blue boxes are formal-based approaches. Black arrows shows how new sub-activities have evolved from existing activities, often as a joint effort.

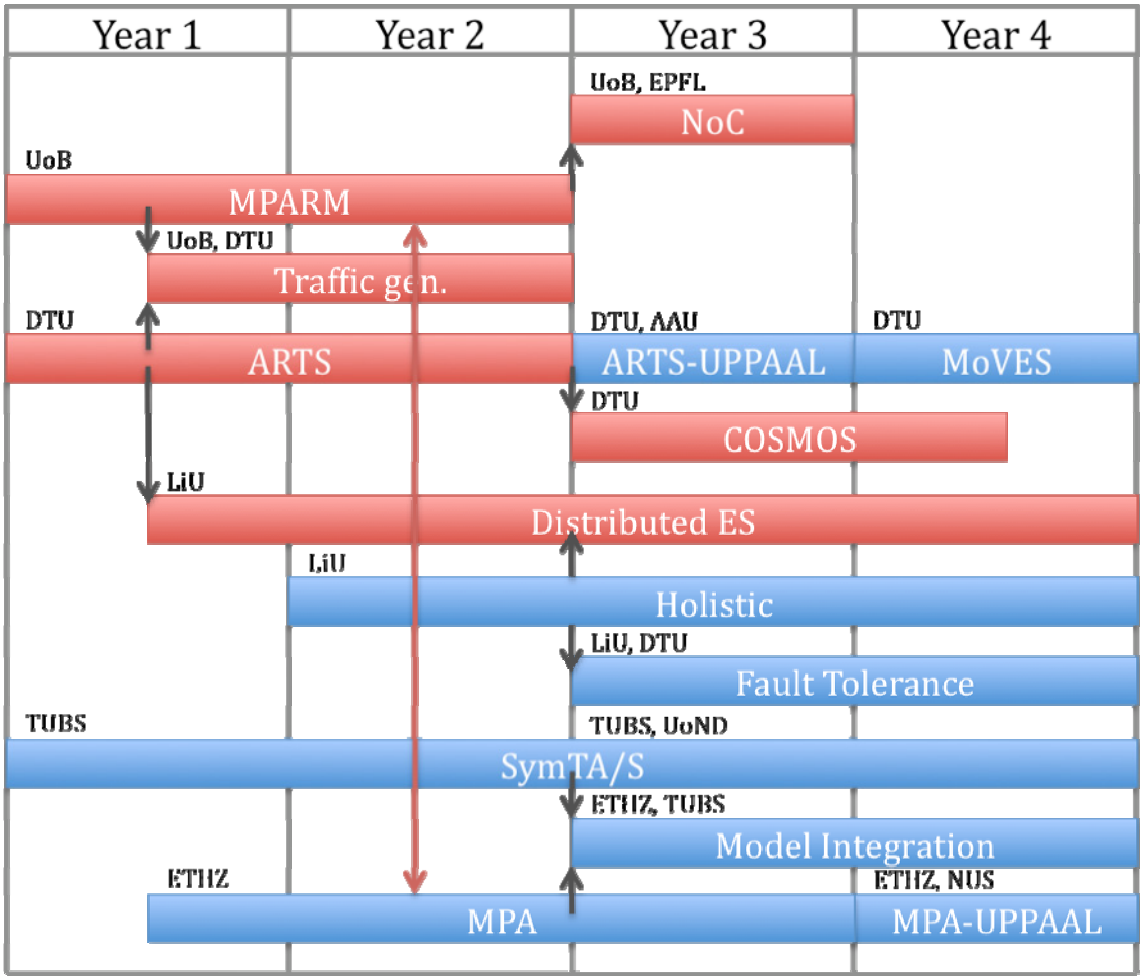

Two trends are clearly seen from the figure; more focus on formal approaches and more interaction towards the end of the project. The increased interaction is also seen from the number of joint publications over the 4 years as shown in the table below.

	Y1	Y2	Y3	Y4
Publications	0	8	8	13
Joint publications	0	2	8	13

TU Braunschweig has provided formal timing models and analysis for Multicore Systems with Shared Memory, or more generally, Multiprocessor-Systems-On-Chip. They have developed methods to integrate the timing implications of shared memory accesses on system performance. They have shown in representative academic examples and industrial case-studies [SHR+08, SNN+08], that these methods allow an accurate performance analysis and optimization during early stages of the design process.

To increase the robustness of real-time systems, TU Braunschweig has provided an analysis methodology of system stability with respect to parameter changes (such as increased execution times). They have provided sensitivity and robustness metrics on the basis of multiprocessor performance analysis, which can then be used in design space exploration and robustness optimization. The framework covers multi-dimensional problems [RHE06], detects scheduling anomalies [RE06], and is largely independent of the utilized performance analysis. The approach can be used to find, on the one hand, system configurations with little design robustness, and on the other hand, to reveal additional performance reserves.

TU Braunschweig has contributed to model coupling in various achievements: The formal performance analysis was accompanied by a power consumption model, developed in cooperation with Prof. Sharon Hu, University of Notre Dame, USA. This combination allows a holistic estimation and optimization with respect to performance and power constraints.

In another achievement, TU Braunschweig has combined the compositional performance analysis with different analysis techniques, mainly a timing analysis of SDF graphs [SSE07], and the real-time calculus from ETHZ [KHET07]. This is possible by relying on a common event model formalism. The resulting analysis can be applied to a larger domain of systems (which could not be captured with one method alone), and to achieve more accurate results by resorting to the more accurate method where possible.

DTU and LiU have investigated fault-tolerant embedded systems. The research proposed has led to several publications, including a best paper award at DATE, and software tools. The application modelling has been performed using a conditional process graph model that can capture the different fault scenarios in the application. Transient faults, which are increasing with hardware integration, have been tolerated using a combination of redundancy techniques, such as replication, re-exection, checkpointing and hardening of processors. The research shows that performance and reliability trade-offs can be supported with optimization tools.

DTU has contributed a simulation-based model, which captures an application executing on a multi-core platform. The model was developed to address the challenges of finding a mapping of the application onto the execution platform, which effectively fulfills the non-functional requirements of the embedded system such as timing, memory usage, energy consumption, and other cost. The simulation model has been extended to a formal model, by expressing it using the semantics of timed automata. The feasibility of the model was shown through a number of academic examples as well as a realistic industrial case.

ETHZ was involved in various activities in the System Modeling Infrastructure cluster. The highlights of this activity can be summarized as follows: Coupling the MPARM simulation framework (Bologna) with an analytical performance evaluation system (ETHZ) in order to combine the best aspects of both approaches. Coupling various performance analysis tools such as Timed Automata (NUS), Symta/S (University Braunschweig) and MPA (ETH Zurich).

Bologna has provided a cycle-accurate simulation framework for modelling multiprocessor systems-on-chip, capable of simulating real-life systems. Throughout the project, this model has been instrumental for many of the related activities. Hence, all partners have been using or interacting with this model. It has been used to extract more accurate high-level models, to allow for simulation of mixed abstraction levels, and to combine simulation- and formal-based models.

3.2 Indicators for Integration

The aim of the activity is to provide a scalable and realistic modelling platform which is abstract enough to provide complete system representation and some form of functional models even for billion-transistor future systems, while at the same time providing the needed flexibility for modelling a number of different embodiments (e.g. multi-processors, homogeneous and heterogeneous, reconfigurable, etc.).

Consistent progress has been reported over the 4 years with respect to the integration indicators. An active and productive cooperation between the partners has been extended towards other clusters and affiliated partners.

From the technical point of view, several new problems have been identified, and will be jointly researched by the partners. The research approach strongly leverages synergies between the partners, by integrating different levels of system abstraction (from scheduling via operating systems to system design). The successful technical cooperation are demonstrated by several presentations and organized meetings and workshops. Besides the many arrangements, the strongest indication of integration is the many tool couplings, integrations and co-developments, as shown in the figure below, as well as the increasing number of joint publications among the core partners.

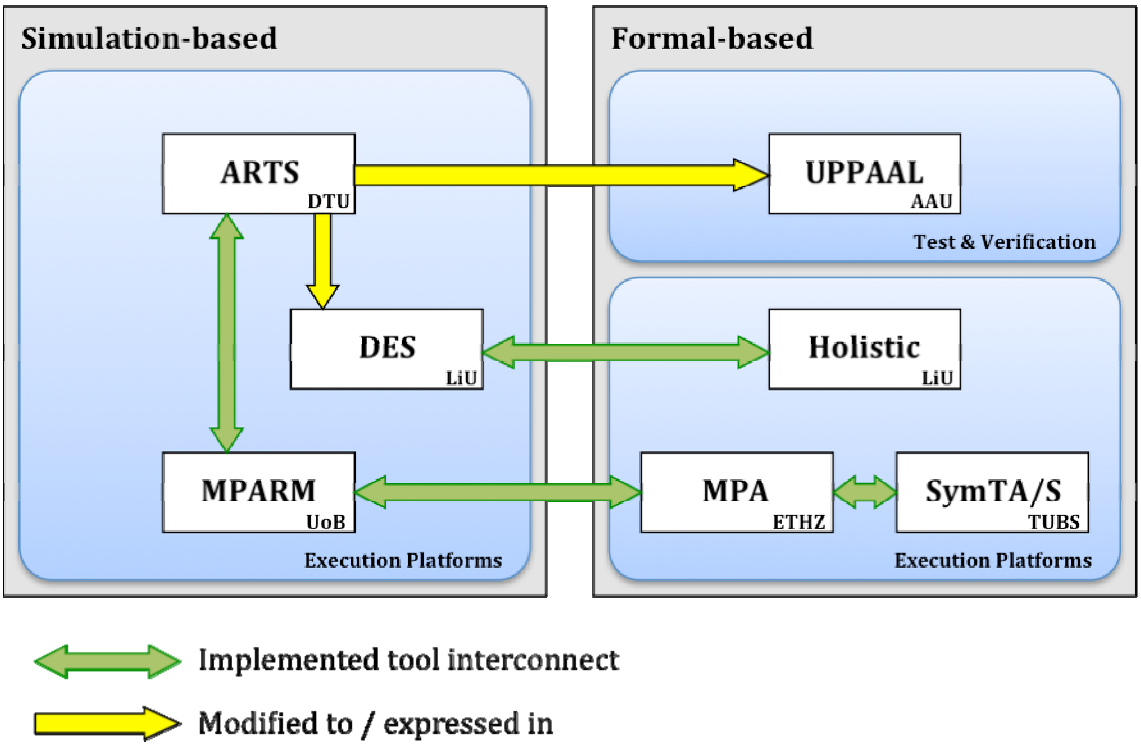

3.3 Main Funding

- Foundation for Strategic Research (SSF), Sweden
- DaNES (Danish Network for Embedded Systems, funded by the Danish Advanced Technology Foundation), Denmark. Period 2007-2010. Budget 9 MEuro, 3.5 Meuro from the Danist Advanced Technology Foundation.

- Hogthrob (project on Sensor Networks funded by the Danish Research Council), Denmark. Period 2004-2007. Budget 0.8 Meuro.
- MoDES (project on Model Driven Development of Intelligent Embedded Systems funded by the Danish Strategic Research Council), Denmark. Period 2006-2009. Budget 0.6 Meuro.
- STMicroelectronics, Italy

3.4 Future Evolution Beyond the Artist2 NoE

The international collaboration on system modelling and performance analysis will be continued in the Artist Design NoE through the research activity on “hardware platforms and multiprocessor system-on-chip”. Many of the activities developed during ARTIST2 will be further developed in this activity. However, a particular challenging topic to be addressed is to obtain a better understanding of the performance analysis an optimization of adaptive systems, i.e. systems that are able to change during run-time in order to optimize metrics such as power consumption, quality of service or reliability and robustness. A specific problem to be addressed is how to performe effective resiurce management in such adaptive systems.

Many of the partners are already participating and collaborating in other EU FP7 projects on methods and tool development, and it is expected that we will see more EU applications also within the ARTEMIS JIT action.

4 Internal Reviewers for this Deliverable

Prof. Paul Pop (Informatic and Mathematical Modeling, Technical University of Denmark)

Prof. Michael R. Hansen (Informatic and Mathematical Modeling, Technical University of Denmark)