



IST-004527 ARTIST2
Network of Excellence
on Embedded Systems Design

Activity Progress Report for Year 3

JPIA-Platform
System Modelling 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	3
1.1	ARTIST Participants and Roles.....	3
1.2	Affiliated Participants and Roles.....	3
1.3	Starting Date, and Expected Ending Date	3
1.4	Baseline	3
1.5	Problem Tackled in Year 3.....	4
1.6	Comments From Year 2 Review	4
1.6.1	<i>Reviewers' Comments</i>	4
1.6.2	<i>How These Have Been Addressed</i>	5
2	Summary of Activity Progress.....	6
2.1	Previous Work in Year 1	6
2.1.1	<i>Simulation-based modeling (month 0 – 6)</i>	6
2.1.2	<i>Formal modeling (month 0 – 6)</i>	7
2.1.3	<i>Simulation-based modeling (month 6 - 12)</i>	7
2.1.4	<i>Formal modeling (month 6 – 12)</i>	8
2.2	Previous Work in Year 2	9
2.3	Current Results.....	12
2.3.1	<i>Technical Achievements</i>	12
2.3.2	<i>Individual Publications Resulting from these Achievements</i>	18
2.3.3	<i>Interaction and Building Excellence between Partners</i>	19
2.3.4	<i>Joint Publications Resulting from these Achievements</i>	20
2.3.5	<i>Keynotes, Workshops, Tutorials</i>	20
3	Future Work and Evolution	24
3.1	Problem to be Tackled over the next 12 months (Sept 2007 – Aug 2008).....	24
3.2	Current and Future Milestones.....	24
3.3	Indicators for Integration	26
3.4	Main Funding	27
4	Internal Reviewers for this Deliverable	28

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 infrastructure

Dr. Magnus Hellring – Volvo (Sweden)

Areas of his team's expertise: requirements analysis.

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 3

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 Year3 has been on:

- Model integration of formal-based models using timed automata models, and extension of current formal-based models to address and encompass more hardware issues, i.e. more issues and details of the execution platform.
- Extending the simulation-based modeling to address issues of dynamically reconfigurable architectures and distributed embedded systems, and to consolidate the integration between the simulation-based models.
- Extending the models to capture issues of fault-tolerance. Addressing simultaneously energy and reliability is especially challenging because lowering the voltage to reduce energy consumption has been shown to exponentially increase the number of transient faults. One aim is to be able to optimize for reliability without degrading energy and latency of the system.
- 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 2 Review

1.6.1 Reviewers' Comments

“General comments apply.”:

All deliverables have now been accepted. To avoid difficulties in the future, it is important that all future deliverables document outcomes achieved through NoE funding relative to the current 18-month plan of work, making a clear distinction between outcomes resulting from NoE funding and outcomes resulting from external funding.

The deliverables were of a uniform excellent quality, written very professionally. The template provides fields for exactly what is needed to report on progress, and the authors have clearly and concisely populated the template in each case. Unlike last year, it is obvious what has been done, in particular, how the NoE funding has supported the integration goals. We have lists of

publications that resulted from the interactions between network members; note that we have not looked at the publications to see if each contains an acknowledgement of ARTIST2 support. We hope this is general practice.

We found the tables describing the primary participants in each cluster in D2 very useful, and feel that the inclusion of a digital photo of the individual is very helpful, not just for the reviewers, but also for anyone outside of the core members who will invariably run into ARTIST2 members at workshops and conferences.

The consortium should put in place a quality process for deliverable. For example, a document from one cluster should be reviewed by independent people from **other** clusters.

The consortium should open itself to external views and additional industries. Today too many stakeholders are left over. We would like to see the number of affiliates growing. Following the recommendations from the last review meeting, we are glad to see that a procedure is in place for this on the website.

The consortium might consider addressing the issues related to fault-tolerance. Today it is treated by verification but not at software level and a real-time system cannot meet its deadline in presence of uncontrolled faults.

Each document should have a short conclusion (what are the results compared to expectation – what specific actions will be taken to enhance things, to get more local funding etc.. (for example)).

External funding figures should be given – it gives an idea of the effort in a particular theme.

Some more graphics in the reports may be welcome (to make reading more pleasant and interesting). PDFs delivered with systematically clickable links in all the reports would be good.

1.6.2 How These Have Been Addressed

No particular comments related to this activity have been given. However, we have tried to follow all general comments given by the reviewers.

In particular, we have started activities in year 3 to address the issues of fault tolerance, which was a general issue raised 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 communication 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 attempts 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 ARTIST2 are the development of an OCP-based interface between computation 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.

Traffic generators (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 Linköping University 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 Linköping groups. A student from DTU has made a short visit to Linköping 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 Linköping is done entirely within ARTIST2. The extensions related to multimedia applications, including the setup and experimentation with 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 express 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 University 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 Current Results

2.3.1 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 3 (covered in the “next 18 months” section from year 2) was to extend the simulation-based modeling to address issues of dynamically reconfigurable architectures, distributed embedded systems, and lab-on-a-chip, and to consolidate the the integration between the simulation models of Bologna and DTU. For the formal-based modeling the focus was on model integration, using timed automata models, and to extend current models to address and encompass more hardware issues.

These objectives have been achieved. For the simulation-based modeling approaches, the ARTS model from DTU has been extended with capabilities for simulating the dynamic behavior of run-time reconfigurable platforms. LiU have finalized their distributed embedded systems simulator derived from the ARTS model, and have used it to validate formal analysis approaches of CAN and FlexRay based systems, by estimating their degree of pessimism. Based on their MPARM simulation framework, Bologna (together with EPFL) have developed a NoC emulation framework which acts as a 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 in order to increase analysis performance. Finally, the model integration of MPARM and ARTS, and the jointly developed reactive traffic generators have been finalized. Within the formal-based modeling approaches, Braunschweig has extended their timing analysis to address shared resources such as memories and busses, issues which are important for modeling MPSoC platforms. Braunschweig have extended their sensitivity analysis to handle multi-dimensional analysis and have proposed expressive robustness metrics for different assumptions and design scenarios, and showed how these can be efficiently considered throughout the whole design process. ETHZ have focused on integrating their MPA modeling formalism with other system models. In particular there has been a deep integration between Symta/S and MPA, which not only entails converters between the two modeling formalisms but also investigations, when to use which formalism. Finally, DTU have expressed most parts of their ARTS model in the timed automata semantics of

UPPAAL and have demonstrated the verification of applications executing on a multi-core execution platform.

Based on the reviewer's recommendation to address issues of fault-tolerance of embedded systems, LiU and DTU have proposed an extension to the process graph model which is able to capture the occurrence of faults. They have shown how design transformations that introduce redundancy, such as re-execution and replication, can be applied on this model. LiU have extended their modeling of NoC platforms to address transient faults, allowing them to optimize NoC communication for fault-tolerance.

In the following, details are given for each sub-activity, listing a title and the partners involved in the sub-activity.

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).

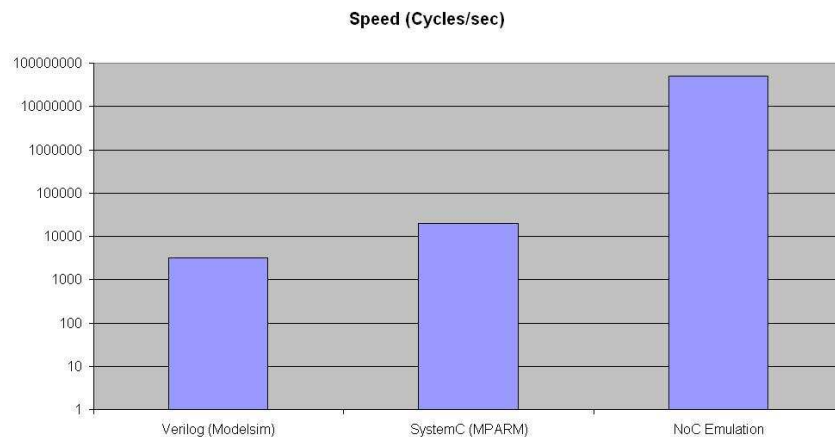


Figure 1: Speed comparison NoC-MPSoC simulation vs.

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-Tolerant 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 Linköping 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 platform and the mapping in a straight forward manner. MOVES then translates the system into a UPPAAL model which is then used to model check the system against given properties. If the model checking fails, the given counterexample produced by UPPAAL, is translated by MOVES into a schedule indicating where the properties were violated. The designer can then use this information to understand why the system failed and to suggest improvements.

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

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

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

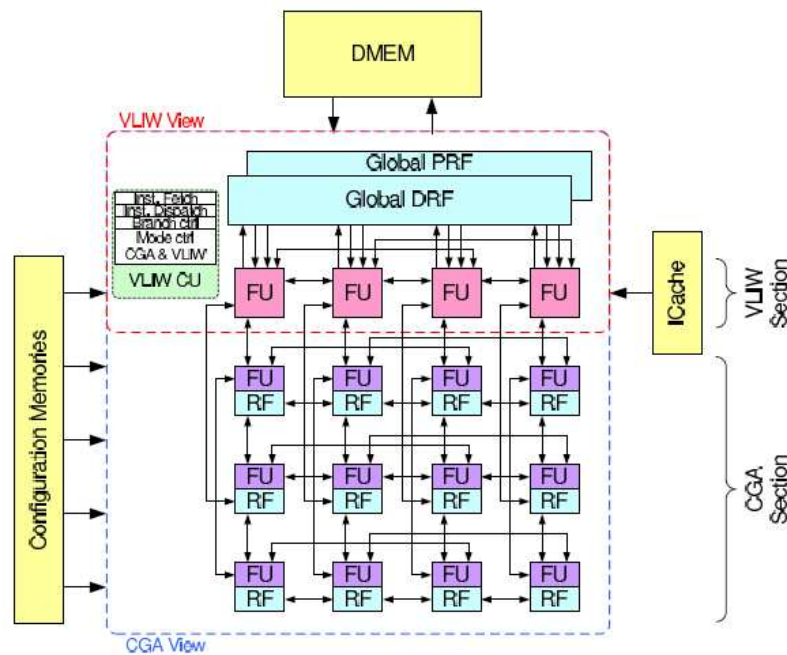
Simulation Platform for Dynamical Reconfigurable Systems (DTU)

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 developed 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 multi-threading. 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 tool domains has been successfully demonstrated. These goals have 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 to 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.

2.3.2 Individual Publications Resulting from these Achievements

TU Braunschweig

[RE06] Razvan Racu and Rolf Ernst. Scheduling Anomaly Detection and Optimization for Distributed Systems with Preemptive Task-Sets. In *12th IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*, San Jose, USA, January 2006.

[HRE07] Arne Hamann and Razvan Racu and Rolf Ernst. Multi-Dimensional Robustness Optimization in Heterogeneous Distributed Embedded Systems. In *Proc. of the 13th IEEE Real-Time and Embedded Technology and Applications Symposium*, January 2007.

[HRE06] Arne Hamann and Razvan Racu and Rolf Ernst. A Formal Approach to Robustness Maximization of Complex Heterogeneous Embedded Systems. In *Proc. of International Conference on Hardware - Software Codesign and System Synthesis (CODES)*, Seoul, Korea, January 2006.

[RHE06] Razvan Racu and Arne Hamann and Rolf Ernst. A Formal Approach to Multi-Dimensional Sensitivity Analysis of Embedded Real-Time Systems. In *Proceedings of the 18th Euromicro Conference on Real-Time Systems (ECRTS)*, Dresden, Germany, January 2006.

Technical University of Denmark

[BHH07] T. Bolander, J. Hansen, M. R. Hansen, Decidability of a hybrid duration calculus, *Electronic Notes in Theoretical Computer Science*, vol. 174(6), pp. 113-133, Elsevier, 2007

[BHM06] Brekling, A., Hansen, M. R., Madsen, J., *A Timed-Automata Semantics for a System-Level MPSoC model*, Nordic Workshop on Programming Theory, 2006

[WuM07] Kehuai Wu and Jan Madsen. *COSMOS: A System-Level Modelling and Simulation Framework for Coprocessor-Coupled Reconfigurable Systems* SAMOS VII: International Symposium on Systems, Architectures, Modelling and Simulation 2007.

Linköping University

[MEP07] Manolache S., Eles P., Peng Z., "Fault-Aware Communication Mapping for NoCs with Guaranteed Latency", International Journal of Parallel Programming, Volume 35, Number 2 / April, 2007.

2.3.3 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:

- ETHZ-Braunschweig: Model and tool integration between MPA from ETHZ and SymTa/S from Braunschweig. The result is published in 1 joint paper.
- DTU-LiU: Proposed an extension to the classical process graph model to handle issues of fault-tolerance. In particular to handle the different fault scenarios that can happen during the occurrence of transient faults. Several visits between DTU and LiU have taken place, and the preliminary results have been documented in 2 joint publications.
- DTU-BOLOGNA: The joint work on traffic generators for network-on-chip simulation has been completed with 3 additional publications.
- DTU-ETHZ-Braunschweig: Participation in the PhD course on Automated Formal Verification of Embedded Systems organized by DTU at DTU, June 4-12, 2007. 2 PhD students from Braunschweig and 3 PhD students from ETHZ participated and gave lectures and hands-on tool demonstrations.

Activities with partners outside the activity

- Bologna-EPFL: Development of a hardware-software NoC emulation framework implemented on an FPGA platform. The framework is used to tune and functional validate on-chip interconnections for MPSoCs. Several visits between Bologna and EPFL have taken place. So far, the main result is the framework itself.
- DTU-IMEC: Extending the ADRES architecture and toolchain to support multitask programming. ADRES is an architecture for coarse grained dynamic reconfiguration. A PhD student from DTU spent 3 month at IMEC working on the multithreaded extension to ADRES. The results of the cooperation are documented in 1 joint publication. A joint journal publication is currently under review.
- 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 and Aske Brekling from DTU visited Oldenburg one week in December 2006, Michael R. Hansen from DTU visited Oldenburg one week in July and Martin Franzle visited DTU one week in August for joint research on formal models for embedded systems. The results are documented in 1 joint publication and presented at 2 workshops.
- DTU-AAU: Modeling the ARTS framework using the timed automata semantics of UPPAAL. Several visits between DTU and AAU have taken place. The results have been presented at 4 workshops, 1 university booth demonstration, and in 1 publication by DTU. A journal publication is currently under review.

2.3.4 Joint Publications Resulting from these Achievements

[KHET07] Simon Künzli, Arne Hamann, Rolf Ernst, Lothar Thiele: Combined Approach to System Level Performance Analysis of Embedded Systems CODES/ISSS 2007, Salzburg, Austria, 2007.

[SSE07] Simon Schliecker and Steffen Stein and Rolf Ernst. Performance Analysis of Complex Systems by Integration of Dataflow Graphs and Compositional Performance Analysis. In *Proc. of Design, Automation and Test in Europe (DATE)*, January 2007.

[SIE06] Simon Schliecker and Matthias Ivers and Rolf Ernst. Integrated Analysis of Communicating Tasks in MPSoCs. In *Proc. 3rd International Conference on Hardware Software Codesign and System Synthesis (CODES)*, Seoul, Korea, January 2006.

[DAM06] P. Del Valle, D. Atienza, I. Magan, J. Flores, E. Perez, J. Mendias, L. Benini, and G. De Micheli, "A Complete Multi-Processor System-on-Chip FPGA-Based Emulation Framework," in *IFIP/IEEE International Conference on Very Large Scale Integration (VLSI-SoC)*, pp. 140-145, 2006.

[CAD06] S. Carta, A. Acquaviva, P. G. Del Valle, M. Pittau, D. Atienza, F. Rincon, G. De Micheli, L. Benini, and J. M. Mendias, "Multi-Processor Operating System Emulation Framework with Thermal Feedback for Systems-on-Chip," in *17th ACM Great Lakes Symposium on VLSI (GLSVLSI)*, pp. 311-316, 2007.

[TVLSI] P. Pop, V. Izosimov, P. Eles, Z. Peng, "Design Optimization of Time- and Cost-Constrained Fault-Tolerant Embedded Systems with Checkpointing and Replication", submitted to *IEEE Transactions on VLSI Systems*, 2007.

[FrH07] M. Fränzle, M. R. Hansen, *Deciding an Interval Logic with Accumulated Durations*, TACAS 2007, LNCS 4424, pp. 201-215, Springer-Verlag, 2007

[MVM06] Kehuai Wu, Andreas Kanstein, Jan Madsen and Mladen Berekovic *MT-ADRES: Multithreading on Coarse-Grained Reconfigurable Architecture* International Workshop on Applied Reconfigurable Computing 2007.

2.3.5 Keynotes, Workshops, Tutorials

Workshop: Workshop on Models and Analysis for Automotive Systems Real-Time Systems Symposium (RTSS)

December 5, 2006

TU Braunschweig participated with a talk on "The Need of a Timing Model for the AUTOSAR Software Standard".

Special Session: Virtual Automotive Platforms Design Automation Conference (DAC)

June 6, 2007

TU Braunschweig was invited to participate in the special session on "Virtual Automotive Platforms" at the renowned Design Automation Conference (DAC). The talk "Automotive Software Integration" showed how formal techniques can be applied to solve performance related integration problems in the design process of complex modern automotive systems.

Tutorial: Extensible Frameworks for System-Level Analysis of Real-Time Systems Real-Time and Embedded Technology and Applications Symposium (RTAS)

April 4, 2006

TU Braunschweig has organized together with ETH Zürich and University of Notre Dame the tutorial "[Extensible Frameworks for System-Level Analysis of Real-Time Systems](#)" at the Real-Time and Embedded Technology and Applications Symposium (RTAS).

Lecture: Supporting Predictable Design Using Formal Analysis Techniques
ARTES Summerschool

August 23, 2006

TU Braunschweig has given a lecture with the title "Supporting Predictable Design Using Formal Analysis Techniques" at the ARTES summerschool (A Network for Real-Time Research and Graduate Education in Sweden) that took place in Nässlingen, Sweden, August 23, 2006. The audience consisted of Phd Students from the field of real-time research, which allowed disseminating recent results in embedded system design to related research teams (mainly) in Scandinavia.

Workshop: Towards a Systematic Approach to Embedded System Design
DATE, Design, Automation, and Test in Europe

Nice, France – 20th April, 2007

DTU has given a talk at with the title "Formalizing the ARTS MPSoC Model in UPPAAL" at the ARTIST2 Workshop at the DATE conference. The target audience of the workshop was industry representatives and researchers wishing to interact about applications and needs for leading-edge Embedded Systems Design tools. The workshop was organized by ARTIST2.

Demo: MOVES, a Tool for Verification of MPSoC Systems
DATE, Design, Automation, and Test in Europe

Nice, France – 20th April, 2007

DTU has given a demo of their tool for modeling and verification of MPSoC systems at the DATE University Booth. In a 2 hour slot, the tool was presented and discussed with academic and industrial people participating in the DATE conference.

Seminar: Quantitative Aspects of Embedded System Design
Dagstuhl seminar

Dagstuhl, Germany – 4-9, Marts, 2007

DTU has given two talks at at the Dagstuhl seminar organized partly by ARTIST2. The two talks were "MOVES: A Tool for Modeling and Verification of Embedded Systems" and "Deciding an Interval Logic with Durations". The purpose of the seminar was to connect the results on performance analysis in ARTIST2 with the community dealing with statistical and stochastic methods. <http://www.dagstuhl.de/de/programm/kalender/semhp/?semnr=2007101>

Workshop: Tool Platforms for Embedded Modelling, Analysis and Validation
CAV, Computer Aided Verification Conference

Berlin, Germany – 1-2 July, 2007

DTU co-organized the ARTIST2 Workshop at CAV 2007. DTU and AAU gave a talk with the title "Validation og Performance Properties with Uppaaal and Applications". The main aim of the workshop was to intensify the cross fertilisation between the formal methods and the embedded systems communities. <http://www.artist-embedded.org/artist/Aims-and-Scope.html>

PhD-course: Automated Formal Methods for Embedded Systems

Lyngby, Denmark – 4-12 June, 2007

DTU has organized an ARTIST2 sponsored PhD course on “Advanced Topics in Embedded Systems”, that took place at IMM, DTU, Lyngby, Denmark, June 4-12, 2007. Lectures were given by ARTIST members from Oldenburg, Germany, ETH Zurich, Switzerland, and Braunschweig, Germany. The course had 17 participants of which 10 were PhD students from 7 different universities spread over 4 different countries. The course were a big success and will be repeated in 2008. <http://www.artist-embedded.org/artist/ARTIST2-PhD-Course-on-Automated.html>

Lecture: Deciding an Interval Logic with Durations
Trust Soft PhD Seminar

Oldenburg, Germany - July, 2007

DTU has given a lecture with the title “Deciding an Interval Logic with Durations” at the Trust Soft Phd Seminar at Oldenburg.

Mini-keynote: Codesign

7th International Forum on Application-Specific Multi-Processor SoC (MPSoC)

Awaji Island, Japan – 25-29 June, 2007

DTU has been given a talk on “*If Fomal Analysis is the Answer – What was the Question?* “. The mini-keynote addressed the problem of verifying complex MPSoC systems using formal methods, in particular addressing the question of which properties of such a system could be formally verified. It presented results from the collaboration between DTU and AAU where the ARTS system has been modeled in Uppaal. <http://tima.imag.fr/mpsoc/>

Workshop: MOVES, a Tool for Modeling and Verification of Embedded Systems
MoDES Workshop

Sønderborg, Danmark – 12-13 Marts, 2007

DTU has given a talk presenting their tool for modeling and verification of embedded systems at the MoDES Workshop held in Sønderborg, 12-13 Marts, 2007.

Summerschool: Advanced Digital Systems Design
Conference name

Lausanne, Switzerland – 25-29th September, 2006

Two members of the cluster on Execution Platforms have been 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

Zurich, Switzerland – 16-17th November, 2006

A Workshop on Models of Computation and Communication brought together scientists from various areas, i.e. formal methods, hardware design and software architecture, see <http://www.artist-embedded.org/artist/MoCC-06.html> .

Workshop: CASTENESS

15.-17th of January 2007

ETH Zurich has been organizing and participating in the CASTENESS Workshop, see www.casteness.org. The workshop 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.

**Workshop: Foundation and Applications of Component-based Design
EMSOFT 2006**

Seoul, Korea, 2006

ETH Zurich has been organizing a Workshop at a major conference in the area of Embedded Software (EMSOFT): “Foundations and Applications of Component-based Design“, October 26th 2006, Seoul. The workshop has been organized in the framework of the Embedded Systems Week (<http://www.esweek.org/>), which federates CODES/ISSS, EmSoft, and CASES.

Dagstuhl Seminar: Quantitative Aspects of Embedded Systems

Sloss Dagstuhl, Germany – 4-9th March, 2007

ETH Zurich has been organizing a Dagstuhl Seminar 04.03.2007-09.03.2007: “Quantitative Aspects of Embedded Systems”. The purpose was to connect the results on performance analysis in ARTIST2 with the community dealing with statistical and stochastic methods. Therefore, organizers of this workshop have been B. Haverkort (Univ. of Twente, NL), J.-P. Katoen (RWTH Aachen, DE), L. Thiele (ETH Zürich, CH), see <http://kathrin.dagstuhl.de/07101/>

Conference: Architecture of Computing Systems (ARCS'07)

Zurich, Switzerland – 12-15th March, 2007

ETH Zurich has been the general chair of the ARTIST2-sponsored conference ARCS'07: “Architecture of Computing Systems”, which took place at the Swiss Federal Institute of Technology (ETH) Zurich, Switzerland, March 12-15, 2007, <http://arcs07.ethz.ch/>. Here, a broad audience was present which allowed disseminating results on embedded system design methods to a larger community.

**Tutorial: Analysis and optimization of real time distributed embedded systems
International Workshop on Embedded Systems**

Seoul October 2006

Petru Eles has given a tutorial at the “International Workshop on Embedded Systems 2006”, Seoul October 2006. With this occasion several results obtained in the ARTIST context have been made accessible to an international audience.

3 Future Work and Evolution

3.1 *Problem to be Tackled over the next 12 months (Sept 2007 – Aug 2008)*

DTU and LiU will continue its work on models for the analysis and optimization of fault-tolerant embedded systems. In particular, they will refine the model to capture fine-grained combinations of several fault-tolerance techniques. At the moment, replication and re-executions can not be combined in the FT-PG model for a single process. There are several ways such a situation can be captured, and they will evaluate the advantages and disadvantages of each approach.

TU Braunschweig will continue its work extending the semantic model of SymTA/S to efficiently cover MPSoC architectures. Additionally, they will further investigate challenging timing issues in multiprocessor systems. Main goals include increasing the precision of latency prediction for multiple memory accesses, and looking into the issue of resource synchronization.

DTU will continue the work on formalizing the ARTS model using timed automata based on UPPAAL. This will allow the same platform model to be expressed as a simulation model and as a formal model. The work will be carried out in cooperation with the research group at CISS in Aalborg which is a partner in the ARTIST2 cluster on test and verification. Work on extending the simulation model of ARTS will be continued.

DTU will refine its formal model to address modeling and verification issues closer to the hardware layer of the execution platform.

ETHZ intends to combine Modular Performance Analysis with timed automata based evaluation methods. This work will be done together with the affiliated partner NUS (National University Singapore). This way, there is a link and integration of MPA with (a) simulation (done in a joint work together with University Bologna), (b) Symta/S (joint work with TU Braunschweig) and (c) timed automata.

LiU will do experimental evaluations using the simulation environment for distributed embedded systems: of particular interest are issues related to fault tolerance.

3.2 *Current and Future Milestones*

In the next 18 months, we will continue our effort in developing and refining the various system models. We will continue our model integration both within and between the simulation-based and formal-based modelling approaches. The model refinement will be based on the feedback gained by using the models, in particular with the activities of Communication Centric Systems and Low Power Design.

TU Braunschweig will continue its work extending the semantic model of SymTA/S to efficiently cover MPSoC architectures. Additionally, TU Braunschweig will conduct further research in sensitivity analysis techniques and its application to predictable system design.

TU Braunschweig have extended their sensitivity analysis to handle multi-dimensional analysis and have proposed expressive robustness metrics for different assumptions and design scenarios, and showed how these can be efficiently considered throughout the whole design process.

TU Braunschweig will continue its work extending the semantic model of SymTA/S to efficiently cover MPSoC architectures. Additionally, they will further investigate challenging timing issues in multiprocessor systems.

University of Linköping will continue its development on the simulation environment based on the ARTS framework from Technical University of Denmark. They will conduct experimental evaluations using the simulation environment for distributed embedded systems. The focus is on evaluation of the impact of various protocols on worst case and average performance; evaluation of pessimism of various response time analysis approaches; impact on quality of control.

LiU have finalized their distributed embedded systems simulator derived from the ARTS model, and have used it to validate formal analysis approaches of CAN and FlexRay based systems, by estimating their degree of pessimism.

LiU will do experimental evaluations using the simulation environment for distributed embedded systems: of particular interest are issues related to fault tolerance. LiU and DTU will continue its work on models for the analysis and optimization of fault-tolerant embedded systems. In particular, they will refine the model to capture fine-grained combinations of several fault-tolerance techniques.

University of Bologna will investigate techniques and approaches to reduce simulation time.

This is essential for analysis of complex platforms and of complex workload, in particular when dealing with cycle-accurate models. To this purpose simulation-acceleration techniques based on hardware emulation will be investigated.

Based on their MPARM simulation framework, Bologna (together with EPFL) have developed a NoC emulation framework which acts as a 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 in order to increase analysis performance. Finally, the model integration of MPARM and ARTS, and the jointly developed reactive traffic generators have been finalized.

Technical University of Denmark will continue research on the ARTS environment.

Extensions will cover; modelling capabilities for dynamically reconfigurable architectures.

This requires that not only the software can be moved and modified during platform execution, but also the hardware itself. The aim is to be able to model and analyse new architectures for reconfigurable computing. The research on modelling wireless sensor networks will be continued. Furthermore, extension of the modelling capabilities toward lab-on-a-chip will be started, in particular towards biochips, i.e. platforms which are able to move microfluidic droplets around within an array of cells in order to mix and analyse chemical liquids.

The ARTS model from DTU has been extended with capabilities for simulating the dynamic behavior of run-time reconfigurable platforms. The extensions towards modeling biochips have not been started. This was mainly due to lack of funding, as a research project together with Duke University in USA did not get funded through the Danish Research Council. The topic is still of interest and DTU is currently searching for funding and have extended the consortium to include people from both physics and biology.

DTU will continue the work on extending the simulation based model towards handling dynamically reconfigurable architectures. In particular, DTU will study different run-time resource management strategies.

Technical University of Denmark will continue the work on linking simulation models with formal models. In particular they will extend their effort in formalizing the ARTS model using timed automaton based on UPPAAL. This will allow the same platform model to be expressed as a

simulation model and as a formal model. The work will be carried out in cooperation with the research group at CISS in Aalborg which is a partner in the cluster on test and verification.

DTU have expressed most parts of their ARTS model in the timed automata semantics of UPPAAL and have demonstrated the verification of applications executing on a multi-core execution platform.

DTU will continue the work on formalizing the ARTS model using timed automata based on UPPAAL. In particular, DTU will refine its formal model to address modeling and verification issues closer to the hardware layer of the execution platform.

ETH Zurich will combine simulation and analytic methods: Continuing the work with University of Bologna and possibility strengthening the cooperating with Technical University of Denmark on combining simulation and analytic estimation methods. The main approach is to use simulation and dedicated benchmark applications in order to profile a hardware platform with respect to OS and communication overhead. Using this information to parameterize analytic performance analysis approaches.

ETH Zurich have focused on integrating their MPA modeling formalism with other system models. In particular there has been a deep integration between Symta/S and MPA, which not only entails converters between the two modeling formalisms but also investigations, when to use which formalism

ETH Zurich will combine Modular Performance Analysis with timed automata based evaluation methods.

3.3 Indicators for Integration

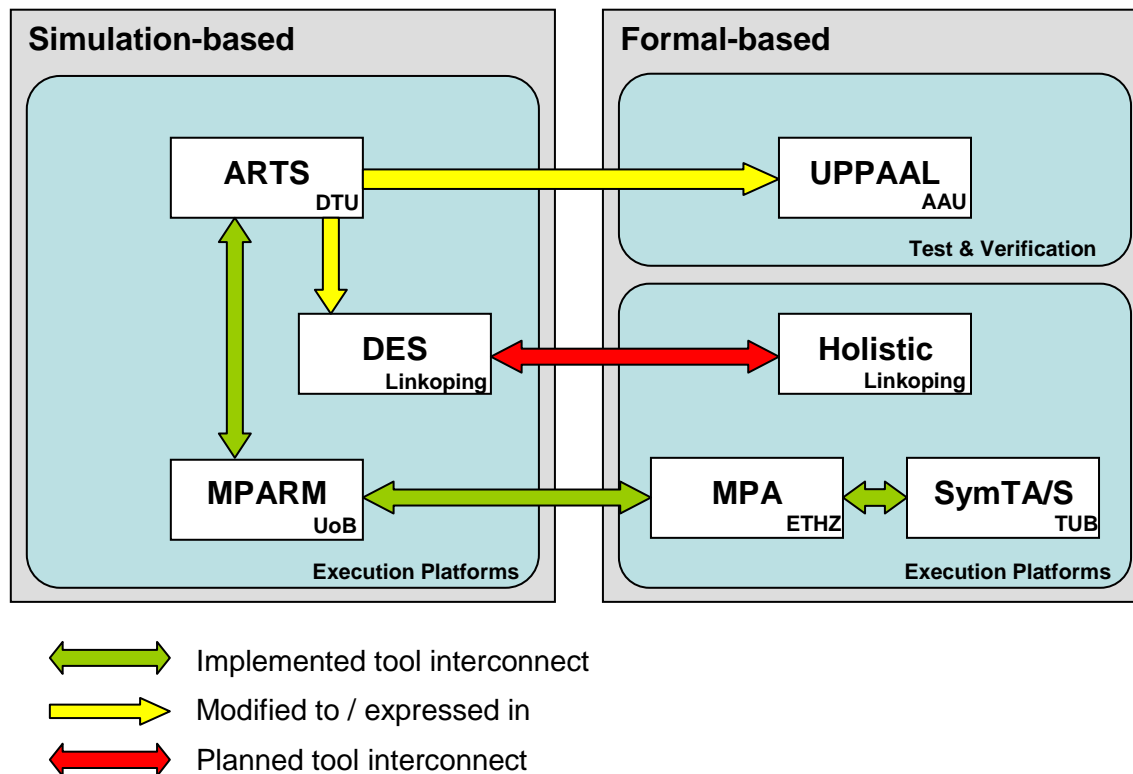
The aim is to provide a scalable and realistic modelling platform which is abstract enough to provide complete system representations and some form of functional models even for billion-transistor future systems, while at the same time providing the needed flexibility for modelling a number of different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

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

For the formal-based models, model and tool integration has been further developed between Braunschweig and ETHZ, and DTU and AAU. This has required meetings and interaction through visits and email, and has resulted in several joint publications and presentations. One outcome has been a PhD course on automated formal verification of embedded systems organized by DTU and with lectures from ETHZ, Braunschweig and Oldenburg. Joint activities on fault-tolerant issues have been initiated among the partners, in particular LiU and DTU have extended classical models to model occurrences of transient faults. This activity has required several visits between LiU and DTU and has resulted in joint publications.

For the simulation-based models, focus has been on finalising model and tool integration and adaptation among the partners. Bologna and DTU has completed their work on traffic generators for NoC simulation, resulting in additional joint publications. Bologna and EPFL have improved simulation performance through the joint development of an emulation framework to functionally validate and tune complex MPSoC systems. IMEC and DTU have worked on models for dynamic reconfigurable architectures which have required longer visits from DTU at IMEC. The cooperation has resulted in joint publications. LiU has completed their simulation model based on the model from DTU and has started to experiment with simulations of automotive application. One aim is to link simulation-based and formal-based models through investigating the degree of pessimism of formal-based approaches as compared to simulation-based.

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 is demonstrated by several joint publications, presentations and organized workshops.



3.4 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

4 Internal Reviewers for this Deliverable

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