IST-004527 ARTIST2
Network of Excellence
on Embedded Systems Design

Activity Progress Report for Year 4

JPRA-Cluster Integration

Communication-centric Systems

Clusters:

Execution Platforms

Activity Leader:

Prof. Dr. Rolf Ernst (TU Braunschweig)

http://www.ida.ing.tu-bs.de/~ernst

The activity assesses the state-of-the-art in models that take into consideration the particularities of various quasi-standard communication protocols during system analysis and scheduling. The communication infrastructure can be optimised in order to be adapted to the particularities of the implemented application. The derived methods meet a growing need in industry to consider formal techniques in embedded system design as a complement to traditional prototyping and simulation based approaches.

Table of Contents

1. Overview of the Activity

1.1 ARTIST Participants and Roles

Prof. Lothar Thiele – TIK, ETH Zürich (Switzerland)
Developing a calculus to describe the performance of communication-centric systems, unifying the models for computation and communication.
Prof. Petru Eles – ESLAB, Linköping University (Sweden)
Schedulability analysis for heterogeneous distributed systems, communication synthesis.
Prof. Rolf Ernst – IDA, TU Braunschweig (Germany)
Formal performance models for networks-on-chip.
Prof. Luca Benini – Micrel Lab, University of Bologna (Italy)
Analytic and simulation based models for performance, power and area of NoCs.
Prof. Jan Madsen – Technical University of Denmark (Denmark)
Power issues in network on chip architectures.

1.2 Affiliated Participants and Roles

Dr. Fabian Wolf – Volkswagen AG (Germany)
Software integration under real-time constraints
Dr. Magnus Hellring – Volvo (Sweden)
Requirements analysis
Dr. Kai Richter – Symtavision (Germany)
Performance analysis of complex distributed systems
Prof. Sharon Hu – University of Notre Dame (USA)
Power analysis and optimization

1.3 Starting Date, and Expected Ending Date

Starting date: September 1st , 2004
Ending date: August 31st , 2008

This is a core activity in ARTIST 2 that was continued throughout the network funding period. With the significant progress and contributions of that activity, the main activity can now be terminated. The results will be included in the ArtistDesign NoE.

1.4 Baseline

Formal communication modelling has been investigated by Zebo Peng and Petru Eles at Linköping University, Lothar Thiele at ETH Zurich (ETHZ) and Rolf Ernst at Braunschweig. The ETH Zurich and Linköping University have outstanding expertise in modelling and analyzing packet flow communication and network processors (ETH) and conditional task graphs combined with statistical modelling. UoB is one of the most widely recognized centres of expertise in NoC design, analysis and road mapping. DTU has a long experience in asynchronous circuits design and in NoC design based on this technology. Embedded architectures and heterogeneous distributed embedded systems have grown to extremely complex computation and communication patterns. New performance models and a corresponding theory are urgently needed. Europe needs to develop skills to safely design such

systems.

1.5 Problem Tackled in Year 4

In year four, we have continued to cover key issues of embedded multiprocessor system design. Joint research has taken place in three major areas: General modelling of real-time systems, multicore architectures, and dependable systems.

Modelling of real-time systems

Communication between tasks in distributed multiprocessor systems often exhibits a complex communication pattern due to the fact that several layers of software (e.g. OS, application) and hardware (e.g. packetization, hierarchical communication protocols, gateways, etc) are involved. It is difficult to capture these patterns with classical event stream models. In the scope of the year 4, we have addressed this issues and investigated hierarchical event stream models, which allow capturing the timing properties of a layered communication in distributed multiprocessor systems. When the hierarchical event stream is transformed, e.g. due to the fact that it passes different system components, the effects on the previously merged streams can be considered separately, which allows the later extraction of the individual event stream information.

System level performance analysis plays a fundamental role in the design process of hard real-time embedded systems. Several different approaches have been presented so far to address the problem of accurate performance analysis of distributed embedded systems in early design stages. The existing formal analysis methods are based on essentially different concepts of abstraction. However, the influence of these different models on the accuracy of the system analysis is widely unknown, as a direct comparison of performance analysis methods has not been considered so far. We defined a set of benchmarks aimed at the evaluation of performance analysis techniques for distributed systems. ETHZ (L. Thiele) with several ARTIST partners applied different analysis methods to the benchmarks and compare the results obtained in terms of accuracy and analysis times, highlighting the specific effects of the various abstractions. The work will be published in the ACM Journal Design Automation for Embedded Systems.

Real-Time Calculus has been used successfully on systems where dependencies between components, via either workload or resource streams, are acyclic. For systems with cyclic dependencies the foundations and performance of the formalism are less well understood. In a joint work with ARTIST partner Bengt Jonsson, Uppsala, ETHZ (Lothar Thiele) developed a general operational semantics underlying the Real-Time Calculus, and use this to show that the behavior of systems with cyclic dependencies can be analyzed by fixedpoint iterations. We characterized conditions under which such fixedpoint iterations give safe results, and also show how precise the results can be. Results were published at EMSOFT08.

For the modelling of heterogeneous distributed real-time systems, different models were proposed in literature and can be used by a system designer. In such a setting, certain components may be modeled in a purely functional manner, while the remaining components require additional modeling of state information. The functional models can be efficiently analyzed but have restricted expressiveness. On the other hand, state-based models are more expressive and offer a richer set of analyzable properties but are computationally more expensive to analyze. Together with affiliated ARTIST partner NUS (P S Thiagarajan), ETHZ (Lothar Thiele) presented a performance analysis technique for distributed real-time systems We showed that by appropriately composing these two classes of models it is possible to leverage on their respective advantages. The resulting modeling technique is as expressive as Event Count Automata, but is amenable to more efficient analysis. We illustrate these

advantages using a number of examples and a detailed case study. The work has been published at RTSS 07.

Multicore and NoC Architectures

The increasing application of multi-core components in real-time systems has required a revised view of the timing hierarchy in formal performance analysis. Due to runtime conflicts on shared resources, such as the memory, the timing of individual tasks can not be determined without knowledge of system timing. We have provided two solutions to tackle this problem:

- LiU has investigated predictable implementation of real-time applications on multiprocessor systems. The emphasis during this last year has been on the issue of efficiency, in the sense of reduced pessimism and minimal overhead. In cooperation with the group at Bologna specific bus controllers for predictable multiprocessor systems have been synthesized.
- An iterative analysis has been proposed by TUBS. This analysis captures the dynamic inter-task interference between cores and allows predicting the worst-case delay also in the case of dynamic run-time bus scheduling. This delay is then integrated into the worst-case response time [SNN+08]. (This analysis is mainly a contribution to Activity D14, System Modelling Infrastructure)

Dependable Systems

The upcoming automotive communication standard FlexRay has for reasons of fault-tolerance two independent channels that can be used for message replication to increase redundancy. However, the protocol specification does not support message acknowledgement. Hence, any fault-tolerance technique that does not involve simple hardware fault-tolerance (transmission on both channels) has to be implemented in the application layer. Several techniques have been proposed for tolerating transient faults: replication on the two channels (the default mechanism), retransmission without acknowledgement in the slots of the static segment and acknowledgement-based retransmission in the dynamic segment. The techniques are based on approaches for the schedulability and optimization of FlexRay, proposed at DTU and LiU, and provide support for trade-offs between timeliness and reliability at the communication level, [Tra08a,Tra08b].

We have also worked on increasing the dependability of autonomous systems with the help of organic algorithms, which provide self-monitoring, self-adaptation, self-repair, and other features known from nature. Future real-time systems, especially in automotive, will exhibit a large space of possible versions and configurations, many of which can not be foreseen at design time, but will be introduced via updates to the customer. To be able to cover system variants in sensitive environments, the system needs to be able to manage itself, monitoring current runtime configurations, accepting or rejecting updates, and perform optimizations based on current resource demands. We have provided a framework that has ported many concepts from established off-line formal performance analysis to the actual target system, where it can be performed in the field [SE08].

1.6 Comments From Year 3 Review

1.6.1 Reviewers' Comments

"General comments apply"

1.6.2 How These Have Been Addressed

Since no particular comments related to this activity have been given, we have tried to follow the general comments given by the reviewers.

2. Summary of Activity Progress

2.1 Previous Work in Year 1

Mixed Performance Analysis in Communication Centric Systems
In a first step, the SymTA/S tool framework that was initially designed to support only the evaluation methods designed at TU Braunschweig was extended with a dynamic library concept, such that different analytical libraries can be loaded into the tool to perform the system-level analysis of embedded systems. In a next step, at ETH Zürich, the formal analysis method Real-Time Calculus was implemented as a Java library that can be used from within the SymTA/S tool. This library was then integrated into the tool and can now be used for performance evaluation.

Hybrid Approach for Performance Analysis of Communication Centric Embedded Systems
After an initial meeting (3 days in Bologna) in March 2004 to discuss the possibilities for a joint effort towards this new approach and exchange the knowledge of the existing performance evaluation methods used, Simon Künzli spent 3 weeks in May 2005 in Bologna for the actual implementation of a case study using such a hybrid approach.
The existing simulation framework was extended by the interfaces needed for the proposed hybrid approach. Further, an example application was analyzed using the new approach. The hybrid analysis can be performed automated and exposes the expected speed-up for the simulation of embedded systems, with only a small deterioration of the accuracy of the results.

Performance Analysis in the System Design Process
In February 2004, Ernesto Wandeler spent 10 days at the ESI. During this time, an appropriate case-study system was identified and analyzed using a formal performance analysis method developed at TIK, ETH Zürich (Real-Time Calculus). Further, Ernesto Wandeler held 3 talks at ESI, to introduce people at ESI to the performance analysis research at TIK. In April 2005,
Marcel Verhoef spent 5 days at ETH. During this time, a journal paper was written, based on a former conference paper. Further, new potential case-study systems, as well as plans for a performance analysis tool were discussed.
As a first case study, an existing distributed in-car radio navigation system was chosen and was specified in UML. For this case study, Real-Time Calculus was used to evaluate and compare 5 different potential system architectures. Sensitivity analysis was applied to all architectures to identify their robustness and potential bottlenecks. For the architecture that is actually used in the commercial implementation of the case study system, the robustness and the bottlenecks could be identified correctly using formal performance analysis methods.

Optimization and analysis of distributed embedded systems
Prof. Eles and his research group, University Linköping, have continued their work in the context of optimization and analysis of distributed embedded systems. They have concentrated on the following issues:

- Analysis of hierarchically scheduled systems
- Timing analysis of distributed task sets communicating through the FlexRay protocol
- Analysis and optimization of distributed embedded systems with fault tolerance requirements

Power analysis and optimization
To initiate the joint activity, Bren Mochocki, University of Notre Dame, spent 2 month at TU Braunschweig. During this time, interfaces between the power analysis tools developed at University of Notre Dame and SymTA/S were created. Based on these interfaces SymTA/S

could be extended with an analysis technique to determine the power consumption of a given embedded system. Since then the power models and the interfaces were regularly extended and refined.

On-chip interconnections for single-chip execution platforms

The work on communication-centric systems by the group of Prof. Madsen, Technical University of Denmark (DTU), has focused on on-chip interconnections for single-chip execution platforms. The starting point for this work has been the development of a clock less NoC architecture (MANGO) and a system-level NoC model based on the multiprocessor simulation environment (ARTS) developed at DTU. The MANGO NoC architecture is based on asynchronous message-passing and provides guaranteed services. Its interface is based on the standard OCP interface protocol which makes the architecture very suited for a modular SoC design flow. The use of a clockless circuit technique has a number of advantages, among which are; inherent global timing closure, low forward latency in pipelines, and zero dynamic idle power consumption. The ARTS modelling framework, which is developed as part of the System Modelling Infrastructure activity of ARTIST2, was extended with capabilities to model interconnect structures. At the system-level, the details of the processing elements and the NoC need to be abstracted in a way that allows for an accurate modelling of the global performance of the system, including the interrelationships among the diverse processors, software processes and physical interfaces and interconnections. To support the designer of single-chip based embedded systems, which includes multi-processor platforms running dedicated RTOS's, with the ability to analyse effects of on-chip interconnect network, the ARTS framework was required to support the analysis of network performance under different traffic and load conditions. This was achieved by extending the model with capabilities for NoC modelling.

Highly scalable communication architectures

The design objective was to develop a NoC targeting heterogeneous systems and featuring support for customizable, domain-specific NoC realizations. This implies designing Xpipes network building blocks as soft cores and to arbitrarily instantiate these blocks so to obtain custom-tailored irregular topologies. All Xpipes components were modelled in SystemC at the cycle accurate level, and integrated in an overall system simulation environment. Network interface was designed with the objective to allow frequency decoupling and efficient protocol conversion between system cores domain and network domain. Moreover, a standard OCP interface protocol with the cores was implemented to increase portability across different platforms. Links design was characterized by the concern to avoid limiting effects of signal propagation time on overall system clock period. This was achieved by providing support for link pipelining and latency-insensitive design. Finally, switch modules design followed the following guidelines: latency minimization, minimum impact of routing logic, output buffering to avoid head-of-line blocking, and initial support for best-effort traffic only. Parameters that can be set at design time include number of switch input/output ports, buffer sizing, flit width, over clocking factor, etc.

High level modelling of system interconnects

High level models for system interconnects were at first derived for state-of-the-art busses, validated on a virtual platform and potentials for their deployment were explored. In particular, cooperation with Linköping University paved the way for exploring different high level models for shared communication resources and for exploiting them within theoretical frameworks for efficient allocation, mapping and scheduling of tasks onto MPSoC hardware platforms.
Specifically, we identified two modelling approaches to on-chip communication (additive models and coarse-grain modelling of communication tasks) and addressed the issue of modelling implicit communication in a predictive way (cache misses, semaphore polling). Then, we developed the theoretical framework taking the Benders Decomposition approach: an

Integer Programming model to assign tasks to computation and storage resources and a Constraint Programming model to schedule tasks onto processors. Non-feasible schedules generate a no-good for the IP problem, which is then re-iterated. The procedure is proved to converge to the optimum.

Hybrid system-level performance analysis approach
High level bus models can also be used for performance estimation. However, it is well known that formal models by themselves may turn out to be inaccurate for exploring the performance of complex multi-processor systems, because abstractions might fail to capture the system behaviour. Similarly, accurate simulation of the entire system might turn out to be infeasible due to the long simulation times. Therefore, we set up cooperation with ETH Zürich with the objective to assess the efficiency of a hybrid approach: combining simulation and formal methods for system-level performance analysis. This approach enables a faster validation of the whole system in that we can decide to model a subcomponent of which the behaviour is well known through a formal analysis, whereas we can have a detailed low level and timeconsuming simulation component modelling for other components. We described how the simulation models can be coupled with the formal analysis framework and showed the applicability of the approach using case studies.

2.2 Previous Work in Year 2

Timing Analysis of the Flexray Protocol
In the second year the Linköping group has continued the work regarding analysis and optimization of distributed embedded real-time systems, with application in automotive electronics. The main goal is to develop models and tools for the analysis and optimization of such communication-intensive systems. Emphasis is placed on the analysis of timing properties, considering the heterogeneous nature of such systems and the particularities of the various communication protocols. In the most recent research the analysis of mixed static/dynamic protocols, such as FlexRay, has been performed [PPE+06].
FlexRay is likely to become a standard for certain automotive applications and the elaboration of the first timing analysis approach for distributed systems built on FlexRay is of importance for our industrial partners. On top of these timing analysis approaches, various system-level optimization tools have been built, performing application mapping, communication synthesis, priority assignment, etc. The Linköping group has closely collaborated with our industrial partners at Volvo as well as with the Braunschweig group. The developed analysis approaches are under integration in the Symta/S environment developed at Braunschweig.

Fault Tolerance
One other issue that has been explored by the Linköping group, in the same context of distributed communication-intensive real-time systems, is that of fault tolerance and, in particular, the issue of transient faults. There are two main aspects of interest here:
(1) Analysis of timing properties in the presence of faults and possible guarantees regarding worst case behaviour
(2) System optimization, such that timing and fault tolerance requirements are satisfied given a certain, limited amount of resources.
An approach for scheduling and worst case analysis with fault tolerance has been developed [IPE+06]. On top of this analysis approach, an optimization technique for task mapping and fault tolerance policy assignment has been elaborated and implemented.

Combination of performance analysis methods: SymTA/S and MPA
This new collaboration is based on collaborations between the two institutions from previous years, where we tried to identify the similarities and differences of the performance evaluation

methods developed (a) at TU Braunschweig integrated in the SymTA/S tool, and (b) at ETH Zurich implemented as toolbox for modular performance analysis (MPA). With this analysis of the weaknesses and strengths of the various methods in mind, we believe that a combination of the methods leads to a significant improvement of analysis results. Especially for systems in which not all parts of the system can be analysed using a single technique due to limitations of the methods, we see the possibility to apply a combined approach which leads to good analysis results. After the analysis of the individual techniques, we are now looking at a common basis for such a combination, and analyse the implementation effort needed for a tool that supports both analysis techniques. The plan for the next months is to implement the changes needed for a combination and analyse an example application to show the strength of the new approach. These steps should also result into a joint publication of the results.
To achieve this, we intend to (1) apply the changes in the tools at each of the partner's sites, (2) organise an integration week where the two parts should be combined to form a single tool, (3) perform the analysis of an example system.

Performance Analysis of an In-Car Radio Navigation System
In this activity, ETHZ investigated an in-car radio navigation system that was specified in UML. Modular Performance Analysis with Real-Time Calculus was used to evaluate and compare 5 different potential system architectures, and sensitivity analysis was applied to all architectures to identify their robustness and potential bottlenecks. For the architecture that is actually used in the commercial implementation of the case study system, the robustness and the bottlenecks could be identified correctly using the above methods. First results on this research were published at the First International Symposium on Leveraging Applications of Formal Methods [WTVL06]. After this symposium, we refined the analysis of the case study system.
Based on the case study system, we also compared a number of different performance analysis and simulation methods. Currently, a hardware test bed is implemented to compare the analysis results with measured results in different system architectures.
The results of the refined analysis, together with a thorough description of the applied analysis methods were published this year in a journal article [WTVL06]. The results of the analysis methods comparison and of the comparison to the measurements will be published in a future joint publication.

Sureal-Project: Hierarchical Event Models
The main goal of the Sureal Project is to define an integrated development process for distributed embedded real-time Systems, especially regarding real-time aspect in all phases of the development. This includes the integration of different techniques for describing, analyzing and modelling real-time aspects. To be able to use different tools specialized in handling realtime aspects in different phases of the system development interfaces must be defined for them to efficiently work together.
Also the early prediction of the timing behaviour, the sensitivity and optimizing possibilities of the architecture play a very important role in such an integrated development process. The tool SymTA/S is capable of analysing such aspects but the underlying methods still have some limitations regarding specific system setups. Up to date, only task sets, which consist of tasks that are activated according to a standard event model can be analysed appropriately. To lift this limitation, first steps towards exploring hierarchical event models are taken. Future Results will be integrated into SymTA/S to further enhance its applicability.

Power Optimization under Timing Constraints
Based on the power analysis extension to SymTA/S which was realized in cooperation with Bren Mochocki during the first project year, TU Braunschweig and University of Notre Dame implemented heuristic and stochastic power optimization algorithms using DVS and SVS (Dynamic/Static Voltage Scaling). The presented algorithms are applicable to complex distributed systems with complex timing constraints (maximum jitter, end-to-end deadlines,

etc.), and are capable of determining Pareto-optimal design trade-offs between system power consumption and timing properties.
The heuristic power optimization approach is based on research of TU Braunschweig related to sensitivity analysis [RHE06], whereas the stochastic algorithms utilize the compositional SymTA/S design space exploration framework [HRJ+06].
The results of this activity lead to a joint publication at the International Conference on Compilers, Architectures, and Synthesis for Embedded Systems (CASES) [RHE+06].

Robustness Optimization for Distributed Embedded Systems
Based on the results achieved in the domain of sensitivity analysis [RHE06], TU Braunschweig developed techniques for optimizing the robustness of embedded real-time systems with respect to variations of system properties like worst-case execution/communication times, bus bandwidth, CPU clock rate, input data rate, etc. Reasons for such variation during the design process or in the field include updates, bug fixes, late feature requests, and product variants.
The developed algorithms consider hard-real time constraints and are capable of optimizing a given system for static and dynamic design robustness. Thereby, the static design robustness optimization approach is applicable to the design scenario where system parameters are fixed early in the design process, whereas dynamic design robustness optimization approach includes possible counteractions to unforeseen system property changes, and is thus applicable to reconfigurable systems.
The results of this research will be published at the International Conference on Hardware - Software Codesign and System Synthesis 2006 (CODES) [HRE06].

Flex Film: High-resolution Real-time Digital Film Applications
In the context of the FlexFilm project, TU Braunschweig developed a multi-board, multi-FPGA hardware/software architecture, for computation intensive, high resolution (2048x2048 pixels), real-time (24 frames per second) digital film processing. The architecture reaches record performance running a complex noise reduction algorithm (used both as example and proof of concept) including a 2.5 dimensions DWT and a full 16x16 motion estimation at 24 fps requiring a total of 203 Gops/s net computing performance and a total of 28 Gbit/s DDRSDRAM frame memory bandwidth. This design was awarded with the "DATE2006 Design Record" distinction [LHR+06].

Simulation-based analysis of SoC interconnection architectures
Industrial MPSoC platforms exhibit increasing communication needs while not yet reverting to revolutionary solutions such as networks-on-chip. The limited scalability of shared busses is being overcome by means of multi-layer communication architectures.
However, the complex interaction among system components and the dependency of macroscopic performance metrics on fine-grain protocol features stress the importance of highly accurate modelling and analysis tools. The work in this area has focused on developing accurate functional model of multi-node on-chip interconnects, as they are currently deployed in high-complexity SoCs today.

Network-on-chip architectures
In the second year the group at the Technical University of Denmark has further developed the NoC architecture called MANGO (*Message-passing Asynchronous Network-on-Chip providing Guaranteed services over OCP interfaces*). In particular the network core, i.e. the routers and links. MANGO is based on clockless circuit techniques, and thus inherently supports a GALS (*Globally Asynchronous Locally Synchronous*) type design flow. This is an advantage in large scale SoC design, since the distribution of a global clock is becoming increasingly difficult.
MANGO employs virtual channels to provide connection-less best-effort routing as well as connection-oriented virtual circuits, for which service guarantees can be given. The

predictability of guaranteed services is a way to promote system-level integrity. The MANGO architecture has been demonstrated through a circuit-level design of a 5x5 router using a 0.13 µm CMOS standard cell library from STMicroelectronics. Netlist simulations showed a performance of 650 Mflits/s under typical timing conditions [BS06]. Three patents [BS05] on the MANGO technology have been filed and a startup company, called Teklatech (www.teklatech.com), was formed as a spin-off from this research. Teklatech is developing a one-step EDA solution to achieving timing closure in large scale, globally synchronous, deep submicron ASIC designs.

Distributed wireless sensor networks
Besides the further development for extending the capabilities of the ARTS system-level modelling framework towards the modelling of wireless sensor networks (reported under the System Modelling Infrastructure action), a sensor node development platform [VLMB05] has been developed, implemented and built. The aim of the platform is to explore hardware/software tradeoffs when designing the node behavior and to calibrate the developed system-level models with real design implementations. In order to efficiently utilize the limited resources available on a sensor node, key design parameters need to be optimized which is only possible by making system-level design decisions about its hardware and software (operating system and applications) architecture.

Simulation-based analysis of SoC interconnection traffic
In Multi-Processor System-on-Chip (MPSoC) design stages, accurate modeling of IP behavior is crucial to analyze interconnect effectiveness. However, parallel development of components may cause IP core models to be still unavailable when tuning communication performance.
Traditionally, synthetic traffic generators have been used to overcome such an issue. 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. Technical
University of Denmark and University of Bologna have jointly developed a reactive traffic generator device [MAMBS05] 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. Even when tested under complex synchronization scenarios, including asynchronous interrupts involving OS interaction in a multiprocessor environment, the proposed traffic generator is able to reproduce IP traffic with full capability to express the application flow.

2.3 Previous Work in Year 3

Comparison of different performance analysis approaches (ETH Zurich)
ETHZ has been leading a major effort in comparing the modeling scope and accuracy of various performance analysis methods: MAST (Univ. Cantabria), Symta/S (TU Braunschweig), Timed Automata and Modular Performance Analysis (ETH Zurich). The results are based on a set of benchmarks that have been determined in a Workshop at Leiden, organized by ETH Zurich. The comparison showed interesting results that will be the basis for future work in this activity. In addition, a joined publication at EMSOFT 2007 [PWT+07] will describe the obtained results in detail. The results have been also presented at various occasions (Workshop on Models of Computation in Zurich, Seminar at EPF Lausanne).
http://www.tik.ee.ethz.ch/~leiden05/index2.html

Mixed performance analysis using MPA and SymTA/S (ETH Zurich, TU Braunschweig)

There has been an intensive cooperation between TU Braunschweig and ETH Zurich in terms of coupling their respective frameworks for performance analysis (Symta/S and MPA). As a

result, one can now mix different model paradigms in a single analysis run and therefore, exploit the different modeling scopes and accuracies. The cooperation has led to a joint publication at CODES 2007 [KHE+07].

Advanced formal analysis features based on evolutionary search strategies (ETH Zurich, TU Braunschweig)

The PISA multi-objective framework from ETH Zurich has been successfully applied to a new system analysis methodology developed at TU Braunschweig: robustness optimization [HRE07a, HRE07b] and system sensitivity analysis [RHE06].

Analysis and optimization of communication-intensive systems (University of Linköping, DTU)

University of Linköping and DTU have continued the work regarding analysis and optimization of distributed embedded real-time systems, with application in automotive electronics. The main goal is to develop models and tools for the analysis and optimization of such communication-intensive systems. Emphasis is placed on the analysis of timing properties, considering the heterogeneous nature of such systems and the particularities of the various communication protocols.
In the most recent research the Technical University of Denmark (DTU) and Linköping University (LiU) have proposed an approach to the timing analysis of applications communicating over a FlexRay bus [PPE+06], taking into consideration the specific aspects of this protocol, including the dynamic segment. More exactly, they have proposed techniques for determining the timing properties of messages transmitted in the static and the dynamic segments of a FlexRay communication cycle. The analysis techniques for messages are integrated in the context of a holistic schedulability analysis algorithm that computes the worst-case response times of all the tasks and messages in the system.
This was the first step towards enabling the use of this protocol in a systematic way for time critical applications. As a second step, DTU and University of Linköping have proposed an optimization approach for determining a FlexRay bus configuration which is adapted to the particular features of an application and guarantees that all time constraints are satisfied [PPE+07b]. Heuristics for solving the bus access optimization problem with FlexRay have been proposed as part of this research. While the optimization techniques proposed by us can also be applied to other heterogeneous distributed applications, solving the particular problem of analysis and optimization of FlexRay-based systems is, today, of particular importance for the automotive industry.

Fault tolerance under energy constraints (University of Linköping, DTU)

One other issue that has been explored by the Linköping group, in cooperation with the group at DTU, is that of fault tolerance and, in particular, the issue of transient faults. During the last year, efforts were concentrated on achieving a certain, required degree of fault tolerance with minimal energy consumption [PPI+07].

Analysis of shared memory accesses in MPSoC architectures (University of Linköping)

The Linkoping group has investigated the impact of memory access over shared buses on the global predictability of real-time systems implemented on multiprocessor SoC architectures. We have elaborated an approach [RAE+07] to system level scheduling and bus access which provides predictable implementations in the context of potential bus conflicts. Bus traffic taken into consideration is not only that for inter-task communication but also that caused by regular

memory access in the case of cache misses. The basic WCET analysis used as part of the proposed systems is based on the Symta/P tool from Braunschweig.

Network-on-chip architectures (DTU)

The Technical University of Denmark (DTU) has continued its research in Network on Chip architectures as described in the following. Work on the MANGO NoC (reported in last year) has addressed the design of the network adapters [BS06b]. One of the experiences from the MANGO NoC is that circuit complexity and power consumption can be quite large. From related work on NoC architectures, it is clear that this is generally the case. For this reason DTU have been studying more "leight-weight" NoC architectures as well as more efficient clock-distribution methods [SBS06], [BSS07]. Finally work has been started on: (1) automatically generating efficient NOC topologies, (2) programming models for NoC-based systems, and (3) combining circuit switching and packet switching in reconfigurable NoC structures.

Simulation-based analysis of SoC interconnection traffic (DTU, University of Bologna)

The Technical University of Denmark and University of Bologna have completed their joint development of a reactive traffic generator device through 3 additional joint publications [MAS07a], [MAS07b], and [MAS07c]. The reactive traffic generator is capable of correctly replicating complex software behaviours in the MPSoC design phase. The developed approach has been validated by showing cycle-accurate reproduction of a previously traced application flow. Even when tested under complex synchronization scenarios, including asynchronous interrupts involving OS interaction in a multiprocessor environment, the proposed traffic generator is able to reproduce IP traffic with full capability to express the application flow.

Fault tolerance and robustness of autonomous systems (TU Braunschweig)

The AIS project ('Autonomous Integrated Systems') is a research cooperation involving several universities. The aim is to research new system design methodologies which enable MpSoCs to realize autonomous behaviour in case of faults and modifications in the environment. This includes self-healing and self-configuration properties on hardware layer due to specially designed components as well as on software layer. The measures on software layer range from an offline design space exploration at early design stages to the adoption of an autonomous operating system, which supports dynamic reconfiguration mechanism at runtime. Our main research interest focuses on the determination of the performance of such systems under real-time conditions. Therefore we consider two generally different types of failures. Transient errors have a bounded duration, and they are countered usually with one-time actions. Examples of transient errors are single event upsets or timing failures due to the difficulty of accurate prediction of transistor behaviour. In contrast permanent errors have a potentially unbounded duration, and mostly they affect complete function units or processors. A defect during manufacturing process causing a processor breakdown as well as a permanent performance reduction due to heat fluctuations are exemplary situations for permanent errors. We implemented algorithms to capture the consequences of this kind of failures and to explore several countermeasures. In this context we developed an approach to analyse the sensitivity of system against performance fluctuations up to complete processor breakdowns. Furthermore possible task migrations have been explored to make the system more robust against errors by remapping functionality within the system in error case. Currently, the system properties during the occurrence of transient errors are investigated, including a derivation of deadline miss probabilities. All analysis algorithms are based on the SymTA/S tool for performance analysis in complex real-time systems. http://www.edacentrum.de/ais

Online performance analysis of distributed embedded systems (TU Braunschweig)

In order to perform online performance analysis in distributed embedded real-time systems, we implemented a framework for distributed analysis of a formal system model [SHE06a, SHE06b]. We based our work on the compositional performance verification approach of Richter et al. as implemented in the tool SymTA/S. A prototype implementation extending the tool has already been tested.
Online Performance control is achieved by using this framework to continuously evaluate the current state of an embedded system. For this, applications entering the system must enter a service contract, stating the computational and communication needs, which are integrated in the current system model. If the resulting system does not violate any constraints, the new application can be accepted into the system. Watchdogs will ensure that no application breaks its service contract. At this time, the concept has been finalized and a prototype has been implemented on a standard PC running parallel processes.
The framework will also be used for optimization of the current system configuration. Heuristic optimization algorithms can perform what-if analyses to continously improve e.g., the robustness of the system. Similar strategies can be applied to deal with faulty hardware, such as degrading processors or if applications break their service contract in order to remain in a feasable system state.

Integrated development process for distributed embedded systems (TU Braunschweig)

Within the SuReal Project, the sureal-process model was defined, which describes the different steps that have to be undertaken for an integrated development process for embedded real-time systems, taking real-time aspects into account. More importantly, it associates the different steps with the different tools defining the interfaces that must be established between them. For the system level timing analysis tool SymTA/S, two interfaces were defined: one between SymTA/S and the UML-tool Ameos from Scopeset and another between SymTA/S and the execution time analyser aiT from AbsInt. The interface between SymTA/S and aiT has already been prototypically implemented using an XML based exchange format. This interface enables SymTA/S to request the execution times of tasks from aiT and incorporate the obtained analysis results from aiT directly into SymTA/S.

Effect of COM-Layers on the communication timing (TU Braunschweig)

In order consider the effects modern COM-Layers have on the communication timing it is necessary to appropriately capture the event stream hierarchies which arise, e.g. when several signals are packed in the same packet. Therefore hierarchical event models, representing these hierarchies, were defined. They not only enable a more accurate description of the timing of the packets transporting the signals, but they also conserve the models describing the inner event streams. This in turn permits to unpack the inner event streams on the receiving side, which can significantly reduce overestimation compared to using non hierarchical event models, e.g. standard event models.

High Speed DDR-SDRAM Memory Controller for the MORPHEUS platform (TU Braunschweig)

Applications designed for the MORPHEUS platform may require a massive amount of memory, as well as sufficient bandwidth, to fully demonstrate MORPHEUS's potential as a high-performance reconfigurable architecture. For example, the film grain noise reduction

application requires massive amounts of bandwidth due to its real-time requirements. To meet these constraints and to eliminate external memory bottlenecks, a high-bandwidth DDR-SDRAM memory controller has been designed for use with the MORPHEUS platform. Using a bandwidth-optimizing two-stage buffered access scheduler, bank interleaving and request bundling, the controller achieves read and write throughput levels up to 2 GiB/s.

Real-time Digital Film Noise Reduction (TU Braunschweig)

A real-time film grain noise reduction algorithm, originally developed for the FlexFilm project to process digital film in real time, has been chosen as an appropriate algorithm to be implemented on the MORPHEUS platform. Film grain reduction is actually itself a combination of different image processing algorithms or tasks, each with massive memory bandwidth requirements.
The following two-phase approach was chosen for implementation. In Phase 1, the film grain noise reduction has been implemented on an existing reconfigurable platform board based on several state of the art FPGAs (FlexFilm board) to closely match the future implementation for the MORPHEUS SoC. Rough estimates show that real time implementation will only be possible for a frame format of about 2 KiPixels x 1.5 KiPixels, which is a widely used digital film format. In Phase 2, which begins in the coming months, the full scale film grain application will be implemented on the MORPHEUS SoC.

2.4 Final Results

2.4.1 Technical Achievements

Communication centric design requires integration of applications running on several communicating platform components. Components are typically heterogeneous and communication networks have multiple stages. In larger systems, even the networks are often heterogeneous with different link types connected over gateways with different scheduling algorithms. Automotive systems are a good example for this type of heterogeneity.

Before Artist 2, only prototyping and incomplete virtual prototyping (simulation) were available for such systems. Formal methods for communication centric design did not scale to larger heterogeneous systems, and simulation only covered small circuits with sufficient details. Other aspects, such as fault resilience, and energy consumption were unexplored. Suitable architectures for multi-core on-chip architectures were just in their infancy.

Artist2 brought major progress to the design, analysis, and optimization of such communication centric systems. In the first 3 years, two existing tools for the modular, scalable analysis of heterogeneous communication centric systems were extended and coupled. These tools, MPA from ETHZ und SymTA/S from TUBS, are based on different formalisms. Coupling required interfaces translating the different event stream semantics. The results showed significant synergies leading to higher analysis accuracy. New tool features were developed including sensitivity analysis and robustness optimization, fault analysis, and communication stack analysis based on hierarchical event stream models.

New architectures, centered around communication were developed and applied, such as the real-time digital film noise reducer that was based on an FPGA platform.

Modeling of Hierarchical Event Streams for Distributed Systems (TU Braunschweig)

TU Braunschweig has continued the work on modeling of COM-layers in distributed real-time systems with the use of hierarchical event models. The basic research was published in [RE08a], and [RE08b], introducing the hierarchical models. This model allows merging several

single event streams to a hierarchical event stream. When the hierarchical event stream is transformed, e.g. due to the fact that it passes different system components, the effects on the previously merged streams can be calculated, which allows the later extraction of the single event streams. These improved models highly improve the accuracy of communication modelling complex distributed real-time systems. The methods have been implemented into the tool SymTA/S.

Online performance analysis of distributed embedded systems (TU Braunschweig)

TU Braunschweig has provided a framework that has ported many concepts from established off-line formal performance analysis to the actual target system, where it can be performed in the field [SE08]. Using this framework, a system is able to manage itself, monitoring current runtime configurations, accepting or rejecting updates, and perform optimizations based on current resource demands, which increases its robustness with respect to late design changes or in-field updates.

Analysis and Optimization of Distributed Embedded Systems & Fault Tolerance (Linköping)

During the four years of the project, the group at Linköping has investigated several aspects related to the optimization and analysis of distributed embedded real-time systems, with application in automotive electronics. The main goal was to develop models and tools for the analysis and optimization of such communication-intensive systems. Several results have been produced and published on the analysis of timing properties, considering the heterogeneous nature of such systems and the particularities of the various communication protocols. On this work the Linköping group has collaborated with the group at TU Braunschweig.

During the second half of the project, in the same context of distributed real-time systems, the research emphasis was on fault tolerance and, in particular, on the issue of transient faults. There are two main aspects of interest here:

(1) Analysis of timing properties in the presence of faults and possible guarantees regarding worst case behavior
(2) System optimization, such that timing and fault tolerance requirements are satisfied given a certain, limited amount of resources.

An approach for scheduling and worst case analysis with fault tolerance has been developed. On top of this analysis approach, an optimization technique for task mapping and fault tolerance policy assignment has been elaborated and implemented.

During the fourth year of the project, fault tolerance aspects of soft real-time systems have been investigated. In this context, the issue of interest is to guarantee the deadlines for the hard processes even in the case of faults, while maximizing the overall utility of the system. We have proposed a novel quasistatic scheduling strategy, where a set of fault-tolerant schedules is synthesized off-line and, at run time, the scheduler will select the right schedule based on the occurrence of faults and the actual execution times of processes, such that hard deadlines are guaranteed and the overall system utility is maximized. The scheduling strategy can also handle overload situations with dropping of soft processes.
This work has been done in cooperation with the DTU group.

Predictable Implementation of Real-Time Applications on MPSoC architectures (Linköping)

As part of its activities in the present project, the Linkoping group has studied the issue of predictable implementation of real-time applications on multiprocessor systems. One of the main difficulties in this context is the problem of memory access due to the shared resources

accessed by parallel tasks. This makes the memory access time (in case of cache misses) very difficult to predict in a way that is not overly pessimistic. The Linköping group has elaborated an approach to system level scheduling and bus access which provides predictable implementations. The basic WCET analysis used as part of the proposed systems is based on the Symta/P tool from Braunschweig.

During the fourth year of the project the emphasis has been on the issue of efficiency, in the sense of reduced pessimism and minimal overhead. An advanced bus access optimization technique has been developed. In cooperation with the group at Bologna specific bus controllers for predictable multiprocessor systems have been synthesized.

Analysis of Distributed Wireless Sensor Networks (ETH Zurich)

Today's wireless sensor networks (WSN) focus on energy-efficiency as the main metric to optimize. However, an increasing number of scenarios where sensor networks are considered for time-critical purposes in application scenarios like intrusion detection, industrial monitoring, or health care systems demands for an explicit support of performance guarantees in WSNs and, thus, in turn for a respective mathematical framework. Based on ARTIST results, we developed a sensor network calculus in order to accommodate a worst-case analysis of WSNs. This sensor network calculus focused on the communication aspect in WSNs, but had not yet a possibility to treat in-network processing in WSNs. As an extension, we now incorporated in-network processing features as they are typical for WSNs by taking into account computational resources on the sensor nodes. Furthermore, we proposed a simple, yet effective priority queue management discipline which achieves a good balance of response times across sensor nodes in the field. The results have been published at RTSS 07.

Generalizing Communication-Centric Performance Analysis (ETH Zurich)

The Modular Performance Analysis based on Real-Time Calculus (MPA-RTC), developed by Thiele et al., is an abstraction for the analysis of component-based real-time systems. The formalism uses an abstract stream model to characterize both workload and availability of computation and communication resources. Components can then be viewed as stream transformers. The Real-Time Calculus has been used successfully on systems where dependencies between components, via either workload or resource streams, are acyclic. For systems with cyclic dependencies the foundations and performance of the formalism are less well understood. In a joint work with ARTIST partner Bengt Jonsson, Uppsala, ETHZ (Lothar Thiele) developed a general operational semantics underlying the Real-Time Calculus, and use this to show that the behavior of systems with cyclic dependencies can be analyzed by fixedpoint iterations. We characterize conditions under which such iterations give safe results, and also show how precise the results can be. Results will be published at EMSOFT08.

FlexRay fault-tolerance aspects (DTU, LiU)

DTU and LiU have investigated fault-tolerance aspects related to the FlexRay protocol. FlexRay has two independent channels that can be used for message replication to increase redundancy. However, the protocol specification does not support message acknowledgement. Hence, any fault-tolerance technique that does not involve simple hardware fault-tolerance (transmission on both channels) has to be implemented in the application layer. Several techniques have been proposed for tolerating transient faults: replication on the two channels (the default mechanism), retransmission without acknowledgement in the slots of the static segment and acknowledgement-based retransmission in the dynamic segment. The techniques are based on approaches for the schedulability and optimization of FlexRay, proposed at DTU and LiU, and provide support for trade-offs between timeliness and reliability

at the communication level, [Tra08a,Tra08b].

NoC arcitectures and programming models (DTU, Bologna)

DTU and Bologna have continued their work on efficient NoC architectures [Mah08]. In particular, DTU has extended their work to include programming models for NoC-based systems, with the aim to address the system designer's point of view. This has been done through a case study of an embedded image processing application [RaStKa08], for which parallelism and scalability has been investigated. The major challenges faced when parallelizing the application were to extract enough parallelism from the application and to reduce load imbalance. The application had limited immediately available parallelism. It was difficult to further extract parallelism since the application had small data sets and parallelization overhead were relatively high. There was also a fair amount of load imbalance, which was made worse by an non-uniform memory latency. Even so, we have shown that with some tuning relative speedups in excess of 9 on a 16-core system can be reached.

2.4.2 Individual Publications Resulting from these Achievements

TU Braunschweig

[RE08a] Jonas Rox and Rolf Ernst. "Modeling Event Stream Hierarchies with Hierarchical Event Models." In Proc. Design, Automation and Test in Europe (DATE 2008), March 2008.

[RE08b] Jonas Rox and Rolf Ernst. "Construction and Deconstruction of Hierarchical Event Streams with Multiple Hierarchical Layers." In Proceedings of the 20th Euromicro Conference on Real-Time Systems (ECRTS), 2008.

[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.

[SR08] Steffen Stein and Rolf Ernst. "Distributed Performance Control in Organic Embedded Systems." In *Autonomic and Trusted Computing (LNCS)*, Volume 5060/2008, pp 331-342, June 2008.

[SHE08] Steffen Stein and Arne Hamann and Rolf Ernst. "Real-time Property Verification in Organic Computing Systems." In *Proc. of the 2nd International Symposium on Leveraging Applications of Formal Methods, Verification and Validation*, November 2008.

Linköping

[AEPR08] Alexandru Andrei, Petru Eles, Zebo Peng, Jakob Rosén, Predictable Implementation of Real-Time Applications on Multiprocessor Systems on Chip, 21st Intl. Conference on VLSI Design, January 4-8, 2008, Hyderabad, India, pp. 103-110.

[RAEP07] Jakob Rosén, Alexandru Andrei, Petru Eles, Zebo Peng, Bus Access Optimization for Predictable Implementation of Real-Time Applications on Multiprocessor Systems-on-Chip, 28th IEEE Real-Time Systems Symposium (RTSS'07), December 3-6, 2007, Tucson, Arizona, USA, pp. 49-60.

ETH Zurich

[JPT+08] Bengt Jonsson, Simon Perathoner, Lothar Thiele, Wang Yi: Cyclic Dependencies in Modular Performance Analysis. ACM & IEEE International Conference on Embedded Software (EMSOFT), ACM Press, Atlanta, Georgia, USA, October, 2008.

[PCT+07] 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.

[SZT07] Jens Schmitt, Frank Zdarsky, Lothar Thiele: A Comprehensive Worst-Case Calculus for Wireless Sensor Networks with In-Network Processing Real-Time Systems Symposium (RTSS 07), IEEE, Tucson, Arizona, US, December, 2007.

[PWT+08] 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.

[CMR+07] 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, 2007.

Denmark Technical University

[RaStKa08] M. Rasmussen, M. Stuart and S. Karlsson “Parallelism and Scalability in an Image Processing Application”, In proceedings of 4th International Workshop on OpenMP, IWOMP’2008, volume 5004 of LNCS, pages 158-169, May 2008.

2.4.3 Interaction and Building Excellence between Partners

Together with Rolf Ernst (TUBS), Lothar Thiele (ETHZ) is co-chairing the dissertation of Arne Hamann (TUBS). During his dissertation, collaborations were conducted on several topics, such as the genetic algorithm framework, model coupling, and quantitative comparisons. These have been the matter of previous reports.

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.

TUBS and ETHZ have collaborated on the coupling of the compositional performance analysis approaches to exploit the different modeling scopes and accuracies. The cooperation has led to a joint publication at CODES 2007 [KHE+07].

Several smaller meetings between ARTIST2 partners took place during the ARTIST2 Summer School 2008 in Europe, September 8-12, 2008, Autrans, France.

Lothar Thiele (ETHZ) has served as faculty opponent for the dissertation of Alexandry Andrei (Linköping), adviser Petru Eles (Linköping). With this occasion several cooperation aspects have been discussed in the area of timing analysis and energy optimisation.

Paul Pop (DTU) has visited Linköping with several occasions during the fourth project year. Cooperation, in particular on the issue of fault tolerant distributed systems, has been discussed during these visits.

2.4.4 Joint Publications Resulting from these Achievements

[KHE+07] Simon Künzli, Arne Hamann, Rolf Ernst, Lothar Thiele. Combined Approach to System Level Performance Analysis of Embedded Systems. In Proc. of the IEEE Conference on Hardware/Software Codesign and System Sythesis (CODES/ISSS), Salzburg, Austria, October 2007.

[Mah08] A Reactive and Cycle-True IP Emulator for MPSoC Exploration Mahadevan, S.; Angiolini, F.; SparsSparso, J.; Benini, L.; Madsen, J. Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on Volume 27, Issue 1, Jan. 2008 Page(s):109 – 122

[PWE+08] S. Perathoner, E. Wandeler, L. Thiele, A. Hamann, S. Schliecker, R. Henia, R. Racu, R. Ernst, M. González Harbour. Influence of Different Abstractions on the Performance Analysis of Distributed Hard Real-Time Systems. In *Design Automation for Embedded Systems*, Springer Science+Business Media, LLC, 2008

[IPEP08] Viacheslav Izosimov, Paul Pop, Petru Eles, Zebo Peng, Synthesis of Flexible Fault-Tolerant Schedules with Preemption for Mixed Soft and Hard Real-Time Systems, 11th EUROMICRO CONFERENCE on DIGITAL SYSTEM DESIGN (DSD 2008), Parma, Italy, September 3-5, 2008.

[EIPP08] Petru Eles, Viacheslav Izosimov, Paul Pop, Zebo Peng, Synthesis of Fault-Tolerant Embedded Systems, Design, Automation, and Test in Europe (DATE 2008), Munich, Germany, March 10-14, 2008, pp. 1117-1122.

[IPEP+08] Viacheslav Izosimov, Paul Pop, Petru Eles, Zebo Peng, Scheduling of Fault-Tolerant Embedded Systems with Soft and Hard Timing Constraints, Design, Automation, and Test in Europe (DATE 2008), Munich, Germany, March 10-14, 2008, 915-920.

[PPIE07] Paul Pop, Kare Poulsen, Viacheslav Izosimov, Petru Eles, Scheduling and Voltage Scaling for Energy/Reliability Trade-offs in Fault-Tolerant Time-Triggered Embedded Systems, 5th Intl. Conf. on Hardware/Software Codesign and System Synthesis (CODES+ISSS), Salzburg, Austria, September 30 - October 5, 2007, pp. 233-238.

[PPEP08] Traian Pop, Paul Pop, Petru Eles, Zebo Peng, Alexandru Andrei, Timing Analysis of the FlexRay Communication Protocol, Real-Time Systems Journal, Volume 39, Numbers 1-3, August, 2008, pp 205-235.

[PPEP+08] Traian Pop, Paul Pop, Petru Eles, Zebo Peng, Analysis and Optimisation of Hierarchically Scheduled Multiprocessor Embedded Systems, Intl. Journal of Parallel Programming, Volume 36, Number 1, February, 2008, pp. 37-67.

[Tra08a] Pop, Traian; Pop, Paul; Eles, Petru; Peng, Zebo, “Analysis and Optimisation of Hierarchically Scheduled Multiprocessor Embedded Systems”, In International Journal of Parallel Programming: Special Issue on Multiprocessor-based Embedded Systems, vol: 36(1), p. 37-67 (2008). Springer Netherlands, 2008.

[Tra08b] Traian Pop, Paul Pop, Petru Eles, Zebo Peng, Alexandru Andrei, “Timing Analysis of the FlexRay Communication Protocol”, In Real-Time Systems Journal, Volume 39, Numbers 1-3, August 2008, pp 205-235.

2.4.5 Keynotes, Workshops, Tutorials

Opening Keynote : Networks, multicore, and systems evolution - facing the timing beast (R. Ernst)
Emerging Technologies and Factory Automation Conference
Hamburg, Germany – September 16, 2008
Embedded systems rapidly grow in several dimensions. They grow in size, from local isolated networks with simple protocols to network hierarchies to large open heterogeneous networks with complex communication behaviour. They grow in performance, from simple microcontrollers to superscalar to multicore systems with many levels of memory hierarchy. And, they expand in the time dimension by moving from static system functions to open and evolutionary functions that change over time and require new design methods and autonomous system functions. All these development contribute to an ever increasing behavioural complexity with equally complex timing. Nevertheless, the fundamental requirements to reliability and performance predictability have stayed and even been enhanced. Embedded system technology has responded with new integration methods and software architectures supported by platform control methods using new service quality metrics, and with composable formal methods that scale with system size. The talk will give an overview on this exiting scientific field and will give practical examples
http://www.etfa2008.org/ETFA_2008/Keynote_Talks.html

Invited Talk : From WCET to system level analysis (R. Ernst)
Workshop on Worst Case Execution Time Analysis
Prague, Czech Republic – June 30, 2008
The goal of the workshop is to bring together people from academia, tool vendors and users in industry that are interested in all aspects of timing analysis for real-time systems. The workshop fosters a highly interactive format with ample time for in-depth discussions. It provides a relaxed forum to present and discuss new ideas, new research directions, and to review current trends in this area.
http://www.artist-embedded.org/artist/Program,1504.html

Invited Talk : ESL Design Education – How much software do we need? (R. Ernst)
European Workshop Microelectronic Education
Budapest, Hungary – May 29, 2008
This talk highlighted challenges for universities to address changing requirements for embedded system designers with an appropriate curriculum. In a growing number of embedded system designs, software development is no more employed as a hardware design companion, subject to an overall design process that starts with a unified system description. It rather comes as a discipline that creates its own layered architecture for a new system quality enabling flexibility, dynamic adaptation, resilience and evolution.
www.eda-publishing.org/ewme2008/htmls/pdfs/I6.pdf

Invited Talk as part of a special session of the DATE Automotive Day : Automotive system level performance analysis and optimization (R. Ernst)
DATE 2008
Munich, Germany – March 12, 2008

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/

Panel : Unifying or overrated – System level design strategy.
DATE 2008, Executive Panel Session
Munich, Germany – March 10, 2008

Invited talk: Synthesis of Fault-Tolerant Embedded Systems (Petru Eles)
DATE 2008 Conference, as part of the special day on Dependable Embedded Systems:
Munich, Germany – March 10, 2008
With this occasion several results obtained in the ARTIST context have been made accessible to an international audience. They are related, in particular, to fault tolerance aspects of distributed real-time systems like those used in automotive applications.

Tutorial: Complex Embedded System Design (L. Thiele).
Lausanne, Sept. 10, 2007.

Tutorial: Analysis of Distributed Embedded Systems (L. Thiele).
CASTNESS WORKSHOP
Rome, Italy, Jan. 15-18, 2008.
http://www.castness.org/twiki/bin/view/CASTNESS/CastNess08

Invited Talk: Workshop Mapping Algorithms onto MPSOC (L. Thiele).
Germany, June 16-17. 2008.

Mini-Keynote : Closing the Loop: Exploration and Estimation (L. Thiele).
MpSoC Conference
St. Gerlach, The Netherlands, June 23-27, 2008.
The talk describes an environment to map applications to MPSoC platforms. The corresponding platform enables the specification, simulation, performance evaluation and mapping of distributed algorithms. Major characteristics are scalability and multi-resolution methods for validation and estimation that combine simulation-based and analytic approaches.
http://www.mpsoc-forum.org/

ARTIST2 South-American School for Embedded Systems 2008 (L. Thiele),
Florianopolis, Brasil, August 25-2, 2008.
After the successful First ARTIST2 South-American School for Embedded Systems in Buenos Aires, Argentina, this second edition in Florianopolis, Brazil, seeks to strengthen the cooperation between Europe and South America in the area of embedded systems, both at educational and research levels. For this purpose, the goal of the school is to provide state-of-the-art courses on embedded systems oriented towards advanced students and young researchers.
http://www.artist-embedded.org/artist/-ARTIST-2-South-American-School-.html

Artitst2 Summer School: Embedded Systems (L. Thiele),
Grenoble, France, Sept 8-12, 2008.
http://www.artist-embedded.org/artist/ARTIST2-Summer-School-2008.html

3. Milestones, and Future Evolution Beyond the NoE

3.1 Milestones

ETH Zurich plans to build on the results of the previous year in terms of comparing different analysis methods in terms of scope and accuracy. As a result, we expect to combine even more different formalisms, e.g. timed automata, and to improve the properties of our own method, Modular Performance Analysis.
This work has been successfully undertaken with several other ARTIST partners. As a result, a conference and a journal publication appeared.

The Linköping group will continue further development of the analysis and optimization techniques for fault-tolerant and predictable distributed systems. In particular, fault tolerance for soft real-time systems will be investigated. The approach to predictable implementation of real-time systems on multiprocessor SoC will be further developed.
Linköping has fully achieved the milestone both related to the fault tolerance aspects and to the predictability issue.

DTU and LiU will continue their collaboration on the FlexRay communication protocol. So far, the topic of transmission reliability has not been addressed during FlexRay optimization. FlexRay has two independent channels that can be used for message replication to increase redundancy. Another option is to retransmit the message in order to protect against faults. DTU and LiU will determine reliability equations for FlexRay messages, and propose an optimization method to increase reliability of transmission. The optimization will decide if a message is to be replicated on the two channels, or retransmitted, such that the timing constraints are still satisfied.
This milestone was fully achieved and have led to two journal publications [Tra08a,Tra08b].

DTU will continue its work on efficient NoC architectures. In particular the will work on techniques to automatically generate efficient NoC topologies. They will investigate reconfigurable NoC structures based on combining circuit switching and packet switching. Finally, work will be started on linking the NoC architecture with the application program through the development of programming models for NoC-based systems.
DTU and Bologna have continued their work on efficient NoC architectures [Mah08]. In particular, DTU has extended their work to include programming models for NoC-based systems, with the aim to address the system designer's point of view. This has been done through a case study of an embedded image processing application [RaStKa08], for which parallelism and scalability has been investigated. The major challenges faced when parallelizing the application were to extract enough parallelism from the application and to reduce load imbalance. The application had limited immediately available parallelism. It was difficult to further extract parallelism since the application had small data sets and parallelization overhead were relatively high. There was also a fair amount of load imbalance, which was made worse by a non-uniform memory latency. Even so, we have shown that with some tuning relative speedups in excess of 9 on a 16-core system can be reached.

TU Braunschweig will further investigate the application of hierarchical event models for performance verification of embedded systems. One promising application will be the detailed analysis of effect that COM-Layers have on communication timing.
TU Braunschweig fully achieved this milestone, Detailed models to consider communication hierarchies have been published in [RE08a] and [RE08b].

Additionally, TU Braunschweig will pursue its research in the domain of online performance verification. Main focus will be the development of heuristic optimization algorithms performing what-if analyses to continuously improve system robustness. Similar strategies will be

investigated to deal with faulty hardware, such as degrading processors or applications breaking their service contract, in order to remain in a feasible system state.
TU Braunschweig has developed and implemented a framework for online performance management on top of embedded hardware and the uC/OS-II real-time operating system. The framework consists of a verification procedure, online monitoring, and reconfiguration protocols. One optimization heuristic built on top of this framework has been proposed in [SHE08], but is not yet implemented.

Additionally, TU Braunschweig wanted to consider timed automata as a further modelling construct to be coupled to MPA and SymTA/S. *TU Braunschweig has not further investigated this topic. ETHZ has however continued to work in this direction.*

3.2 Indicators for Integration

As a result, we have developed an important prerequisite for the design of embedded systems, namely analytic methods to estimate system properties. In particular, the approaches followed so far in the context of Real-Time Systems and Embedded System Design Communities are integrated and related to each other.

3.3 Main Funding

- Hogthrob (project on Sensor Networks funded by the Danish Research Council);
- Intel Corp. (university research grant (Compose project));
- Bosch (direct industrial grant),
- SRC (coordinated project with Penn State University);
- Swedish Foundation for Strategic Research (SSF);
- Centre for Industrial Information Technology (Linköping University);
- IBM Research Rueschlikon (Network Processor Design).
- DFG (Deutsche Forschungsgemeinschaft)
- BMBF (German Bundesministerium für Bildung und Forschung)

3.4 Future Evolution Beyond the Artist2 NoE

ARTIST2 has inspired work on several topics that could not be fully pursued within its lifetime, but have sparked research into directions that now continue within other projects.

For example, the integration of different levels of Quality-of-Service will be continued to be investigated in the AIS Project (http://www.edacentrum.de/ais/). In particular, mixing real-time and non-real time processing in a multi-core systems is the focus of the Compose project (http://www.ida.ing.tu-bs.de/en/research/projects/compose/)

The challenges to cover autonomous systems, with self-adaptive functions known from organic computing has become a Schwerpunktprogramm of the DFG (http://www.organic-computing.de/spp).

Several partners of this activity will also continue their fruitful collaboration in the scope of the ArtistDesign the COMBEST projects.

4. Internal Reviewers for this Deliverable

- Bengt Johnsson (Uppsalla University)