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

Transversal Activity Progress Report for Year 2

Transversal Activity:
Design for Predictability and Performance

Transversal Activity Leader:
Prof. Bengt Jonsson (Uppsala)
http://user.it.uu.se/~bengt/

Policy Objective (abstract)

Embedded systems are required to satisfy requirements on predictability of timing, memory, processing power, power consumption, etc. They also have increasing demands on (average) performance. The objective of this activity is to develop technology and design techniques for achieving predictability of systems built on modern platforms, and to investigate the trade-offs between performance and predictability. This work will need to be carried out in a synergistic manner, involving all levels of abstraction in embedded systems design, spanning from high-level requirements to detailed implementation details on specific platforms.

Versions

number	comment	date
1.0	First version delivered to the reviewers	December 18th 2009

Table of Contents

1. Overview

1.1 High-Level Objectives

Embedded systems in many application domains are required to satisfy strict requirements on timing, while respecting limited supply of resources in terms of memory, processing power, power consumption, etc. All systems also have increasing demands on (average) performance, which has motivated the introduction of efficiencyincreasing features which drastically increase variability and decrease predictability and analyzability. Since the introduction of new architectural features is inevitable, it is important to

- develop technology and design techniques for achieving predictability of systems built on modern platforms, and
- investigate the trade-offs between performance and predictability.

This work will need to be carried out in a synergistic manner, involving all levels of abstraction in embedded systems design, spanning from high-level requirements to detailed implementation details on specific platforms, and is therefore the subject of a transversal activity involving all clusters of the NoE.

-- Changes wrt Y1 deliverable --

No changes with respect to Year 1.

1.2 Industrial Sectors

Predictability is an important system requirement in all sectors of embedded systems, whose operation should not fail for different reasons. An obvious sector is that of safety-critical systems, which arise in transportation, power automation, medical systems, and related areas. The market for safety-critical embedded systems is large and steadily increasing. According to a study by the international ARC Advisory Group with headquarters based in Dedham, Massachusetts, the safety systems and critical control system market, which was around $650 million in 2003, will grow at an average annual rate of over 7 percent per year to over $900 million in 2008. ARC's Safety and Critical Control System Worldwide Outlook Market Analysis and Forecast Through 2008 predicts a healthy growth of the safety system market for process industries over the next five years.

The industry developing safety-critical embedded systems is severely suffering from design practices leading to unpredictable system behaviour. The determination of guarantees for non-functional requirements is postponed to a late design stage, and then often fails because of design decisions taken earlier. Establishing a methodology reconciling predictability and efficiency will have a very strong impact on systems-design and implementation practice in industry.

Predictability is important also in other sectors, where systems failure may lead to economic consequences, as in consumer electronics, telecom, etc.

-- Changes wrt Y1 deliverable --

No changes with respect to Year 1.

1.3 Main Research Trends

Predictability can be regarded as an effect of choosing suitable hardware and software architectures (in a wide sense) that lead to systems whose worst-case behaviour is easy to predict, and of utilizing analysis techniques that are able to provide these guarantees for the chosen system architecture. Important architectural considerations occur on many levels in a system hierarchy. As a general rule, static allocation of resources leads to predictable systems, whereas dynamic allocation makes predictability difficult. Challenges addressed by this activity appear at all levels of abstraction in the design process

- **Modeling and Validation of systems and of components**: Principles and structures for system and component modeling that are conducive to achieving predictability, by allowing *a priori* predictability analysis and by allowing mappings to platform architectures that preserve predictability. Investigations of how modeling and analysis techniques extend to non-traditional system structures, including distributed and networked architectures, for which predictability is more difficult to achieve. Precise definitions and characterizations of the central concepts, including *predictability, robustness.* Exploring trade-offs between predictability, resource consumption, and performance.
- **Compiler Techniques and Program Analysis**: Timing analysis, i.e., predicting the worst-case execution time (WCET) of a piece of code, is a hard problem, but significant breakthroughs have been obtained in recent years for many types of processors. Commercial tools, all from Europe, are available. Research in timing analysis is closely dependent on research on system-design concepts that increase predictability The issues stretch from the processor architecture across all layers to the application and is caused by the variability of execution times. The goal is to increase the predictability of system behaviour. An important issues is also timing analysis for compilation, especially in the light of multiple processors and other architectural features. An important goal is to marry timing analysis with compilation, in order to make timing properties immediately visible to the embedded systems developer.
- **OS/MW/Networks**: On the operating system level, scheduling and reservation of resources is a widely researched topic, with a vast literature. Operating system mechanisms, such as scheduling, mutual exclusion, interrupt handling and communication, can heavily affect task execution behaviour and hence the timing predictability of a system. For example, preemptive scheduling reduces program locality in the cache, increasing the worst-case execution time of tasks compared with non-preemptive execution. The object-oriented programming style, although attractive as a software development methodology, introduces dynamics into the execution time by the dynamic binding of methods to calls. Techniques that improve predictability include schemes that a priori reserve resources in a wide sense. This can be in the form of reserving time slots for execution of tasks, reserving time slots for communication between tasks (e.g., in the time-triggered architecture and in the synchronous programming paradigm). In future research, it is important to explore the tradeoff between performance and predictability in scheduling. Also important is to investigate of software architectures for time-predictable real-time operating systems, with the goal to avoid that the execution of OS code adversely affects the time-predictability of application tasks and vice versa, thus making the computation-time needs of both operating system activities and application tasks easily predictable.
- **System and Processor Architecture:** Simple processor architectures lead to more predictable systems than complicated ones. Current architectures include many features that decrease predictability, such as implicit concurrency, e.g., pipelining, super-scalarity, out-of-order execution, and dynamically scheduled multi-threading. The restricted processor-memory channel-bandwidth and the growing speed gap between

processor and memory has led to the introduction of deep memory hierarchies and several types of speculation. Dynamic power management technology, which is critical for reducing the power consumption of hardware, also has a significant impact on predictability. Research on predictability has considered, e.g., to replace dynamic memory management by static and predictable ones, such as scratchpads, to characterize and develop more predictable replacement policies in dynamic caches,

The current introduction of multicore processors provides new challenges to predictability, since they introduce new concurrency and communication needs to system development. It is not yet clear how to build predictable and performant systems on multicore platforms.

-- Changes wrt Y1 deliverable --

Polishing 1st and 2nd bullets.

2. State of the Integration in Europe

2.1 Brief State of the Art

The problem of WCET determination has been solved for single tasks and several types of processors and some replacement strategies, including least-recently-used (LRU). Higher degrees of predictability in the cache system can be achieved by taking decisions statically instead of dynamically. Compiler-directed memory management using scratchpad memory, originally developed to decrease energy consumption, also increases time predictability. Reactive processors are also promising because they allow the direct predictable execution of synchronous languages (Esterel), thanks to the direct support of the multi-threading and of the synchronization between threads. Analysis of scheduling policies has been well-researched for single processor systems, but is still not a resolve task for multicore platforms. From the hardware point of view, system interconnects present a significant challenge to predictability, in that they are shared among multiple communication actors (cores, IOs, accelerators, etc.). Time-triggered communication protocols have been proposed, among others, to enhance interconnect robustness and predictability. Techniques for general analysis of timing and researches in predictable, often distributed, embedded systems model messages and communication resources in a similar way as tasks and computation resources. They start from a restricted event model, e.g. periodic, sporadic or periodic with jitter, and have been able to provide analysis results where the interference between event triggered and time triggered computation and/or communication paradigms can be bounded. A unifying approach to performance analysis has beeen proposed based on real-time calculus.

-- Changes wrt Y1 deliverable --

No changes with respect to Year 1.

2.2 Main Aims for Integration and Building Excellence through ArtistDesign

Predictability is a concern which cuts vertically across levels of abstraction in embedded systems design, spanning from high-level requirements to detailed implementation details on specific platforms. It therefore needs to be carried out in a synergistic manner, and is therefore the subject of a transversal activity involving all clusters of the NoE. Previous activities in ARTIST2 have primarily focussed on integration for different layers of abstraction: hardware platforms, compiler technology, timing analysis, modelling, etc. The main purpose of this activity is to integrate research teams working on differen levels of abstraction in embedded systems design.

-- Changes wrt Y1 deliverable --

No changes with respect to Year 1.

2.3 Other Research Teams

Much of the cutting-edge research is performed in Europe, to a large extent by Artist Design Partners.

The group at Univ. of Saarland with its spin-off company AbsInt is world-wide leading in the area of timing analysis of hard real-time systems. The development of aiT, the timing-analysis tool of AbsInt, is based on many years of research on static analysis. Static analysis of an embedded program is used to derive invariants about execution states for all inputs to the program. These invariants allow the derivation of reliable upper and lower bounds on the execution times of programs on a given hardware architecture. The ETHZ group has been developing analytic methods based on max+ algebra to analyze combined computation and communication systems, alloiwing a modular approach to analysis of performance and predictability of distributed hardware-software systems. TU Dortmund is a leader on the combination of compiler and architectural techniques for predictable embedded systems. The Uppsala team has developed UPPAAL, a leading model checker for analyzing timed systems. The Univ. of Bologna has produced several significant contributions in the area of low power design, power management and energy-predictable system design. The group has also pioneered the concept of network-on-chip, a new paradigm for building scalable and efficient on-chip communication fabrics for next-generation multi-core platforms. The Vienna team has developed leading architectures and protocols for predictable networked systems. York is one of the leading groups concerning techniques for designing real-time predictable systems.

UC Berkeley (California) and Columbia University (New-York) have jointly developed a time-predictive architecture, called PRET, based on a shared six-stages pipelined non-speculative processor, a scratchpad memory, and thread-level parallelism. Each individual thread executes at a relatively slow, but very predictable rate.

The two teams worldwide that are leader in the design of reactive processors are the University of Kiel (Germany) and the University of Auckland (New Zealand). Reactive processors are dedicated to execute very efficiently reactive programs such as Esterel. The instruction set offers direct support for preemption, suspension, synchronization between threads and with the environment, and so on. Examples of reactive processors include KEP from the University of Kiel, and EMPEROR (singlethreaded version) and StarPRO (multithreaded version) from the University of Auckland.

-- Changes wrt Y1 deliverable --

Slightly updated, two paragraphs added at the end.

2.4 Interaction and Building Excellence between Partners

During the operation of the ARTIST2 FP6 network of excellence, links have been developed between groups working on compiler techniques for achieving predictability of code. Examples include the cooperation between ETH Zurich (Lothar Thiele) and Saarland University (Reinhard Wilhelm), working on timing analysis, and the TU Dortmund (Peter Marwedel), working on compiler techniques. These links are used, e.g., to develop techniques for “predictability-aware compilers”. In particular, TU Dortmund is working on the integration of timing analysis and compilation. The corresponding work on the WCC-Compiler led to a general interest in this type of work. TU Dortmund also works on resource-aware compilation. The main focus is on memory-architecture aware compilation, implemented through pre-pass compilation tools. This work also includes the mapping to multi-processors. The work is performed in cooperation with ICD and the MNEMEE-project at ICD.

As another example, IST Austria (previously EFPL), University of Salzburg and PARADES are collaborating on compositional languages and approach to embedded system that can be considered the extension of the Giotto approach and are reminiscent of the Metropolis and Ptolemy work carried out at Berkeley.

Several partners are collaborating on the development of bus access policies in multiprocessor designs, with the goal to design predictable, yet efficient, designs. Concerned partners include UoB, ETHZ, SSSA Pisa, Linköping, Braunschweig, TU Vienna,:

- Several phone meetings/discussions have occurred during 2009.
- The enhanced MPARM platform from Bologna has been extensively used by ETHZ and SSSA Pisa and other partners. This has led to frequent interactions between the groups.

Several partners are collaborating on characterizations of basic concepts related to predictability, including, .e.g., robustness. These issues were discussed, e.g., at the RePP Workshop (ESWEEK, Grenoble, Oct. 2009).

-- Changes wrt Y1 deliverable --

The old text (1st paragraph) has been revised, and following paragraphs added.

2.5 Interaction of the Transversal Activity with Other Communities

The predictability activity interacts with several ongoing European projects. These include Predator, which aims at developing a research and design discipline that looks at predictability and efficiency in a synergistic maner, involving all levels of abstraction and implementation in embedded system design. Several ArtistDesign Partners are active in Predator.

Bologna, Dortmund, and Vienna participate in the HIPEAC NoE: this will help establishing links between ArtistDesign and the computer architecture and compiler community, through presentation in HIPEAC-organized events (e.g., workshops, ACACES summer school).

The Dortmund group promotes education in embedded systems through a published text book ("Embedded System Design"). Several other groups were asked to comment on upcoming second edition of the book, to be published in spring 2010. The group also organizes the WESE workshop on embedded system education. The group leader teaches at ALARI (Lugano) and is the European editor for the new Springer series on embedded systems (see http://www.springer.com/series/8563).

TU Dortmund organizes the SCOPES series of workshops on compilation for embedded systems, the workshop on the mapping of applications to MPSoCs and the workshop on software synthesis, held during the embedded systems week 2009.

Saarland University is participating in the German nationally funded project, Automatic Verification and Analysis of Complex Systems (AVACS), which is among others concerned with validating timing-analysis methods and tools as well as timing properties of embedded systems.

The COSTA (Compiler-Support for Timing Analysis) project at TU Vienna, funded by the Austrian Science Fund, focuses on techniques for compilers to support WCET analysis. One of the main goals within the project is to make code more predictable, where the elimination of timing anomalies by appropriate code generation strategies (i.e., without the need to changing the hardware) is a central aim within the project.

-- Changes wrt Y1 deliverable --

Has been updated to reflect current situation.

3. Summary of Activity Progress

3.1 Technical Achievements

The technical work involves all levels of abstraction in embedded systems design, spanning from high-level requirements to detailed implementation details on specific platforms. Let us consider them in some order.

Modeling and Validation of systems and of components

Robustness and Predictability of Embedded Systems (IST Austria)
Robustness and predictability were identified as main challenges in embedded system design and were considered from a conceptual viewpoint in work by IST Austria [Hen08], reported in Deliverable 15-(7.2)-Y1. The technical contribution of [Hen08] was to suggest how predictability can be formalized as a form of determinism, and robustness as a form of continuity. IST Austria continued the effort on studying robustness and predictability in embedded systems in Year 2, following several directions:

- We provided a novel definition of robustness for sequential circuits, based on a notion of continuity. We characterized the exact class of circuits that are robust with respect to our definition and developed an algorithm for checking robustness of arbitrary circuits [DHLN09].
- We developed a method for synthesizing robust systems from high-level specifications [BGHJ09].
- Hierarchical Timing Language (HTL) is a real-time coordination language for distributed control systems. We developed techniques for modular checking of HTL program properties that guarantee that the program exhibits predictable behavior [HKMS09].
- In the context of synthesis of predictable real-time systems, we provided a translation from real-time high-level specification languages Metric Temporal Logic (MTL) and Metric Interval Temporal Logic (MITL) into deterministic timed automata [NP09].

Methodology for designing component-based real-time systems. Universidad de Cantabria
During the last year, work has continued on the development of a methodology for designing component-based real-time systems. We have defined a strategy that can be followed by the designer of a real-time component-based application to configure its schedulability, satisfying the opaqueness requirement that is typical of the components paradigm. This strategy relies on RT-D&C, an extension of the D&C specification that incorporates metadata about the temporal behaviour of components and platforms, which allow the designers of the applications to analyse their temporal behaviour, and to extract from this analysis the schedulability configuration parameters that must be assigned to the component instances and to the platform in order to guarantee the fulfillment of the timing requirements of the application. These parameters are incorporated to the deployment plan of the application. The applications are executed on top of the RT-CCM technology. Managing only the information included in the deployment plan, the launching tools can configure the set of container services that have been defined in the technology for managing the schedulability of the applications in an opaque way, without accessing the code of the components.

Participation in the standardization of the MARTE UML profile for embedded real-time systems (Universidad de Cantabria)
Efforts have continued in standardization and dissemination of the results achieved with the publication of the MARTE standard, in particular in the scope of the OMG, primarily for the finalization of MARTE but also participating in the RTF of SysML and the definition of a new mapping for IDL to ADA2005. We have sent and solved a number of issues in the MARTE standard as members of the finalization task force and enrolled in the new revision task force. We have also participated in the MARTE USERS Group, and defined modeling strategies for the usage of MARTE in the design of avionics applications using ARINC653 standard platforms.

As part of the ADAMS EU project, wee have organized the STANDRTS'09 workshop as satellite event of ECRTS for giving visibility to the standardization efforts made in the real-time and embedded systems communities.

Withing the EVOLVE project we have delevoped initial proposal for a modeling methodology to address early validation and verification, in particular for doing scheduling analysis with MAST from UML-MARTE models.

Timing analysis and Compiler Techniques

Timing Analysis and Timing Predictability (USaar and AbsInt)
We have improved the techniques for timing analysis in the presence of several microarchitectural features in several ways.

Whereas classical research work often used LRU replacement policy to achieve good predictability, FIFO replacement is wide-spread in practice. We have invented a generic policy-independent framework for cache analysis that couples may- and must-analyses by means of domain cooperation. With this framework we achieved a precise may-analysis for caches with FIFO replacements, increasing both the predicted hits and misses.

Regarding pipeline behaviour, WCET analysis has to consider many possible hardware states. Classical representations of the pipeline domain enumerate all reachable states explicitly. For modern, complex pipelines, the analysis can become infeasible due to memory and computation time constraints. To combat this problem, we have devised a symbolic representation of pipeline states and their transitions as binary decision diagrams. This compact representation leads to significant savings in memory consumption and execution time during WCET analysis.

On the software side, we worked on the derivation and exploitation of operating modes in embedded software. We identified heuristics to derive operating mode candidates from source code and a procedure to exploit mode information to arrive at mode-specific WCET bounds. Furthermore, we developed two methods to deal with the additional uncertainty of cache effects arising from dynamic memory allocation. In a first approach, we are developing a predictable dynamic memory allocator. In a second approach, we developed a tool to automatically transform dynamic into static allocation with comparable memory consumption. For both approaches, the necessary information about the dynamically allocated memory is derived by an adapted shape analysis together with appropriate abstraction techniques.

WCET Analysis in the Presence of Context Switches (USaar, AbsInt, SSSA)
In preemptive systems, not only the WCET, but also the context switch costs need to be bounded. In case of preemption, cache memories may suffer interferences between memory accesses of the preempted and of the preempting task. These interferences lead to some additional reloads that are referred to as cache-related preemption delay (CRPD). This CRPD constitutes the major part of the context switch costs. The original concept to bound the CRPD uses the concept of useful cache blocks (UCB). These are memory blocks that may be in cache before a program point and may be reused after it. In recent work, we tightened the

CRPD bound by using a modified notion of UCB: Only cache blocks that are definitely cached are considered useful by our approach. Such blocks are called DC-UCBs. The computed CRPD based on DC-UCBs (when used in combination with the bound on the WCET) delivers a safe bound on the execution time in case of preemption. Experimental results show that our approach provides up to 90% tighter CRPD bounds. A prototype implementation has been integrated into the aiT Timing Analyzer by AbsInt and shipped to Pisa for further collaboration on the analysis of the intertask timing behaviour.

Integration of timing analysis and compilation (TU Dortmund, AbsInt)
The work on Dortmund's Worst-Case Execution Time-aware C Compiler WCC has been continued in ArtistDesign Year 2.

Related to the compiler's core infrastructure, a sophisticated fully automatic loop analyzer based on polyhedral models has been added to WCC. It is able to determine loop iteration counts and flow facts required for WCET analysis for broad classes of loops. Furthermore, a back-annotation module is able to map WCET-related timing data from low-level assembly code back to high-level C code such that a WCET timing model now is available a C code level.

On top of the WCC infrastructure, several compiler optimizations known from the standard literature have been made WCET-aware. This includes Function Inlining, Loop Unswitching and Loop Unrolling. Since the impact of many of these optimizations depends on complex characteristics of the entire application currently under compilation, and of the underlying hardware, machine learning techniques steering these optimizations have been integrated into WCC.

Finally, WCC is currently extended to support code generation and optimization for multiple tasks of multi-process applications. In this area, a multi-task cache partitioning optimization has been developed.

Overall, the WCC compiler can be considered the leading WCET-aware compiler. The integrated tool set allows studying the impact of optimizations for WCET minimization. This achievement concerns item 2 of section 1.3.

http://ls12-www.cs.tu-dortmund.de/research/activities/wcc

Enforcing Time-Predictability on Contemporary Computer Architectures (TU Vienna)
Variable timing of instruction execution and memory accesses as well as the dynamic allocation of resources in contemporary ES computer architectures cause so-called timing anomalies, i.e., observations of local timing phenomena are not consistent with global observations (e.g., when running the same program from two different hardware states, one might observe a local execution-time decrease but a global execution-time increase when comparing the first and the second run). We started to explore strategies for eliminating timing anomalies by different code-generation techniques: Inserting nop instructions resp. instructions that create new register-use dependencies in the code removes any choices for dynamic resource allocation and thus anomalous timing effects. The cost of this technique is an increase in code size, though the observed growth in code size was smaller than expected. This work on the software-based elimination of timing anomalies will be continued in year 3.

OS/MW/Networks

Integrating Scheduling Analysis and Model Checking (Uppsala, York)
One important element of predicatbility is the results obained by undertaking schedulability analysis of the system. For single processor systems such analysis is mature for both fixed priority abd EDF dispatching. For multiprocessor platforms there are still a number of open issues that means that optimal performance (of the hardware) and tight means of prediction are still not available - although they are the subject of considerable research both within

ArtistDesign and elsewhere. These activities and resulting publications are reported on in the Scheduling and Respource Management deliverable (D10-(5.2)-Y2) and are therefore not repeated here.

Partitioning the shared caches on multicores for timing predictability (Uppsala)
The major obstacle to use multicores for real-time applications is that we may not predict and provide any guarantee on real-time properties due to the on-chip shared resources such as L2 cache. In this work, we propose to use cache space isolation techniques to avoid cache contention for hard realtime tasks running on multicores with shared caches. We have presented a scheduling strategy for real-time tasks with both timing and cache spaceconstraints, which allows each task to use a fixed number of cache partitions, and makes sure that at any time a cache partition is occupied by at most one running task. In this way, the cache spaces of tasks are isolated at run-time. We have developed a sufficient schedulability test for non-preemptive fixed-priority scheduling for multicores with shared L2 cache, encoded as a linear programming problem. Our experiments show that the test which employs an LP solver can easily handle task sets with thousands of tasks in minutes using a desktop computer.

Timing Predictability on Multi-Processor Systems with Shared Resources (ETHZ)
Multi-processor systems are becoming increasingly important in consumer electronics as well as in industrial applications, such as automotive software. Tasks need to share data across processing unit boundaries, e.g., local variables, triggering the need for a communication fabric. Real-time tasks execute periodically on a processing element mand are constituted by sequential superblocks.

We consider several models to schedule the superblocks and organize accesses to the shared resources within the superblocks. First, superblocks can be executed sequentially, i.e., a superblock is activated as soon as its preceding superblock has finished, or they can be executed according to a static schedule (preassigned time slots). Second, we consider three models to access shared resources: (1) dedicated access model, in which accesses happen only at the beginning and the end of a superblock, (2) general access model, in which accesses could happen anytime during the execution of a superblock, and (3) hybrid access model, which combines the dedicated and general access models. We show the relation between these models with respect to schedulability and provide experimental results that reveals the model that performs best.

Architecture and System Design

Predictable, Fault-tolerant Embedded Systems Design: Analysis and Optimization of Fault-Tolerant Embedded Systems with Hardened Processors (Linköping, DTU)
Linköping University and DTU have an ongoing collaboration concerning the design of embedded systems which have to behave in a predictable way even in the presence of transient faults. During the last year the emphasis of the work has been on the analysis and optimisation of fault-tolerant hard real-time embedded systems, based on an approach in which hardware and software fault tolerance techniques are combined. The basic trade-off is between processor hardening in hardware and process re-execution in software which, together, have to provide the required levels of fault tolerance against transient faults with the lowest-possible system costs. We have proposed a system failure probability (SFP) analysis that connects the hardening level with the maximum number of re-executions in software. We have developed design optimization heuristics, to select the fault-tolerant architecture and decide process mapping such that the system cost is minimized, deadlines are satisfied, and the reliability requirements are fulfilled.

Prof. Paul Pop from DTU has visited Linköping on several occasions during this period.

Predictability for Multiprocessor SoC Architectures (Linköping, Bologna, Braunschweig, TU Vienna)

One of the major issues in the context of predictability for multiprocessor systems is the shared communication infrastructure. The traffic on the bus does not solely originate from data transfers due to data dependencies between tasks, but is also affected by memory transfers as result of cache misses. A bus access policy and bus access schedule has to be developed which (1) guarantees predictability and (2) provides efficiency in terms of system performance. We have developed an overall strategy and framework for predictable multiprocessor applications. We have addressed the issues of bus access optimization and bus controller design. Bus access optimization is crucial in achieving predictability while, at the same time, maintaining efficiency in terms of performance. In order to demonstrate the practicality of the approach, we have designed and synthesised adequate bus controllers fro the proposed protocols.

During Year 2, our main emphasis has been on analysing and reducing the hardware and software overhead required by guaranteeing predictability. We have worked on analysing the trade-off between average and worst case performance and on developing optimisation approaches which produce an implementation that delivers a good average performance without compromising worst case behaviour and keeping, at the same time, predictability.
The Symta/P tool for WCET analysis, from Braunschweig and the MPARM framework from Bologna have been used in this project.

Complementing the work at Linköping on synchronous time-slided memory-bus arbitration, TU Vienna explored this static bus arbitration paradigm for tasks that are not synchronized to each other and the bus arbiter. Besides demonstrating the time-predictability of this execution scheme, an iterative algorithm for the optimization of the bus-access schedule has been developed. This algorithm “steals” memory bandwidth from tasks that meet their deadlines easily and redistributes this bandwidth to those tasks that do not (yet) meet their timing constraints. Further, it could be shown that the use of single-path code within the tasks makes it possible to obtain constant task execution times on chip-multiprocessor systems.

New bus model for MPSoC’s system bus analysis and optimization (University of Bologna)

A new model of the communication architecture has been implemented. The Transactional Level Model bus module (TLM BUS) takes into account features of modern high-performance communication buses, namely the capability of supporting burst interleaving, multiple outstanding transactions and split transfers. The bus is packet-based, i.e., a ”Transaction” on the interconnect is composed by several packets.

One of the main aspects UoB focused on is the impact of different interconnection arbitration policies on overall system performance. In Real Time systems bus arbitration policy is one of the key aspects for achieving the best trade-off between performance and predictability. Several arbitration policies can be implemented on it. We developed different arbiters, namely:

- Round Robin. It gives equal bandwidth to all masters preventing their starvation, but gives no predictability guarantees,
- Pure TDMA policy. Implicitly ensures a good predictability degree, but bus performance is strongly dependant on slots allocation: sub-optimal slot assignment may result in a waste of bandwidth,
- Two Layers TDMA + Round Robin policy. It is a mix of the previous ones and has been implemented in modern buses.

UoB are now focusing on performance evaluation of the different arbitration policies regarding both the speed and the predictability for the entire system architecture.

Development of power prediction algorithms for energy harvester (University of Bologna)

Several wireless sensor network and distributed embedded system applications leverage

energy harvesting technologies such as small size photovoltaic modules. The advantage of solar energy over other forms of environmental energy is that the available solar power can be predicted with reasonable accuracy allowing the implementation of efficient power management techniques. However accurate predictions of future energy profiles can be expensive in term of memory occupancy and complexity and a trade-off between accuracy and computational effort must be evaluated. UoB developed different solar energy prediction algorithms that give estimates of future available energy over the time. They are computationally simple and have a small memory footprint to facilitate the implementation in resource limited solar powered embedded systems.

A time-predictive programming language and architecture (INRIA and CAU Kiel)
In collaboration with the University of Auckland, INRIA and the University of Kiel have designed a new time predictive programming language, called PRET-C. It is an extension of C with reactive constructs inspired from the Esterel synchronous programming language. These extensions allow the user to program parallel threads with a notion of logical synchronization barrier and with reactive inputs and outputs for communicating with the environment. Communication between the parallel threads is achieved with shared C variables. At the same time, because PRET-C is an extension of C, complex data types can be handled much more easily than with Esterel. Thanks to the synchronous inspired extensions, the semantics of PRET-C is deterministics. Also, since the extensions are defined as C macros, we rely on the gcc macro expander for compiling into C code. Time predictability is then achieved thanks to a WCET algorithm that allows us to compute very precisely the worst-case reaction time of the PRET-C program, thanks to a translation of the program into a timed UPPAAL model. We have also designed a dedicated execution architecture, called ARPRET, based on a reactive processor coupled with a general purpose soft-core processor. This architecture is multi-threaded. It offers native support for the reactive constructs of Esterel-like programming languages: this distinctive feature makes it time-predictable without sacrificing throughput. Accordingly, PRET-C can be either compiled into C or into the ARPRET assembly code. This offers both time predictability and excellent performance, compared to comparable programming languages (Esterel, SyncCharts, or C with proto-threads). During year 3, we plan to work with UC Berkeley on implementing a common benchmark suite and on memory hierarchy issues.

-- The above is new material, not present in the Y1 deliverable --

3.2 Individual Publications Resulting from these Achievements

University of Bologna

[BBB09] C. Bergonzini, D. Brunelli, and L. Benini, Algorithms for Harvested Energy Prediction in Batteryless Wireless Sensor Networks, in: 3rd IEEE International Workshop on Advances in Sensors and Interfaces, pages 144-149, 2009

Universidad de Cantabria

[LDM09] Patricia López Martínez, José M. Drake, and Julio L. Medina. "*Enabling Model-Driven Schedulability Analysis in the Development of Distributed Component-Based Real-Time Applications*". 35th Euromicro Conference on Software Engineering and Advanced Applications, Component-based Software Engineering Track, Patras, Greece, August 2009, IEEE, ISBN 978-0-7695-3784-9, pp. 109-112.

[LCD09] Patricia López, César Cuevas y José M. Drake. *"RT-D&C: A real-time extension of the OMG's Deployment and Configuration Specification"*. Workshop on the Definition, evaluation, and exploitation of modelling and computing standards for Real-Time Embedded Systems

(Standarts 2009), held in conjunction with the 21st Euromicro Conference on Real-Time Systems (ECRTS 09). Dublin, Irlanda, Junio 2009.

[PVPM09] Pablo Peñil, Eugenio Villar, Héctor Posadas, and Julio Medina. *"SystemC executable specification of the MARTE generic concurrent and communication resources under different Models of Computation"*. In Proc. of Workshop on the Definition, evaluation, and exploitation of modelling and computing standards for Real-Time Embedded Systems. STANDRTS'09. Satellite workshop of the the 21st EUROMICRO Conference on Real Time Systems,. Dublin, Ireland, July 1-3, 2009,

TU Dortmund

[LGMM09] P. Lokuciejewski, F. Gedikli, P. Marwedel and K. Morik. *Automatic WCET Reduction by Machine Learning Based Heuristics for Function Inlining*. Proceedings of SMART '09: 3rd Workshop on Statistical and Machine Learning Approaches to Architectures and Compilation, January 2009, pp. 1-15.

[LCFM09] P. Lokuciejewski, D. Cordes, H. Falk and P. Marwedel. *A Fast and Precise Static Loop Analysis based on Abstract Interpretation, Program Slicing and Polytope Models*. Proceedings of CGO '09: International Symposium on Code Generation and Optimization, March 2009, pp. 136-146.

[LGM09] P. Lokuciejewski, F. Gedikli and P. Marwedel. *Accelerating WCET-driven Optimizations by the Invariant Path – a Case Study of Loop Unswitching*. Proceeding of SCOPES '09: 12th International Workshop on Software & Compilers for Embedded Systems, April 2009, pp. 11-20.

[PLM09] S. Plazar, P. Lokuciejewski and P. Marwedel. *WCET-aware Software Based Cache Partitioning for Multi-Task Real-Time Systems*. Proceedings of WCET '09: 9th International Workshop on Worst-Case Execution Time Analysis, June 2009, pp. 78-88.

[LoMa09] P. Lokuciejewski and P. Marwedel. *Combining Worst-Case Timing Models, Loop Unrolling, and Static Loop Analysis for WCET Minimization*. Proceedings of ECRTS '09: 21st Euromicro Conference on Real-Time Systems, July 2009, pp. 35-44.

ETHZ

R. Pellizzoni, A. Schranzhofer, J.-J. Chen, M. Caccamo, L. Thiele, Worst Case Delay Analysis for Memory Interference in Multicore Systems, accepted for Publication, DATE 2010, March, Dresden

Andreas Schranzhofer, Jian-Jia Chen, Lothar Thiele, "Timing Predictability on Multi-Processor Systems with Shared Resources," in Workshop on Reconciling Performance with Predictability (RePP), in conjunction with ESWEEK (Embedded Systems Week), Oct. 11-16, 2009, Grenoble, France.

IST Austria

[NP09] D. Nickovic, N. Piterman. From MTL to Deterministic Timed Automata, Technical Report 2009/2 Imperial College London, 2009.

USAAR

[HR09] J. Herter and J. Reineke. *Making dynamic memory allocation static to support WCET analyses*. Proceedings of the 9th International Workshop on Worst-Case Execution Time (WCET) Analysis, June 2009.

[AB09] S. Altmeyer and C. Burguière. *A new notion of useful cache block to improve the bounds of cache-related preemption delay*. Proceedings of the 21st Euromicro Conference on Real-Time Systems (ECRTS), July 2009.

[GR09] D. Grund and J. Reineke. *Abstract interpretation of FIFO replacement*. In Proceedings of the 16th International Symposium on Static Analysis (SAS), August 2009

[LPW09] P. Lucas, O. Parshin, and R. Wilhelm. *Operating mode specific WCET analysis*. Proceedings of the 3rd Junior Researcher Workshop on Real-Time Computing (JRWRTC), October 2009.

Uppsala University

N. Guan, M. Stigge, W. Yi and G. Yu. *New Response Time Bounds for Fixed Priority Multiprocessor Scheduling.* Proc. RTSS09, 30th IEEE Real-Time Systems Symposium, December 1 - 4, 2009 Washington, D.C., USA.

N. Guan, M. Stigge, W. Yi, G. Yu: *Cache-aware scheduling and analysis for multicores*. ACM & IEEE International Conference on Embedded Software (EMSOFT), ACM Press, Grenoble, France, 2009, pp 245-254.

TU Vienna

M. Schoeberl and P. Puschner. *Is Chip-Multiprocessing the End of Real-Time Scheduling?* In Proc. 9th Euromicro Workshop on WCET Analysis, p. 96-106, 2009.

M. Schoeberl, P. Puschner, and R. Kirner. *A Single-Path Chip-Multiprocessor System*. In Software Technologies for Embedded and Ubiquitous Systems, 7th IFIP WG 10.2 International Workshop, Springer, LNCS, 2009.

-- The above are new references, not present in the Y1 deliverable --

3.3 Joint Publications Resulting from these Achievements

[ARGT09] S. Andalam, P.S. Roop, A. Girault and C. Traulsen, PRET-C: a new language for programming precision timed architectures (extended abstract). Proceedings of RePP'09: International Workshop on Reconciling Performance with Predictability, Grenoble, France, October 2009.

[BGHJ09] R. Bloem, K. Greimel, T. Henzinger, B. Jobstmann. Synthesizing Robust Systems, In *Formal Methods in Computer Aided Design (FMCAD'09), 2009*

[DHLN09] L. Doyen, T. Henzinger, A. Legay, D. Nickovic. Robustness of Sequential Circuits, submitted for publication, 2009.

[GRG09] D. Grund, J. Reineke, and G. Gebhard. *Branch target buffers: WCET analysis and timing predictability*. In Proceedings of the 15th International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA), August 2009.

[HKMS09] T. Henzinger, C. M. Kirsch, E. R. B. Marques, A. Sokolova. Distributed, Modular HTL, in Real-Time Systems Symposium (RTSS'09), 2009.

[IPPEZ09] V. Izosimov, I. Polian, P. Pop, P. Eles, and Z. Peng. Analysis and Optimization of Fault-Tolerant Embedded Systems with Hardened Processors. Proceedings of DATE: Design Automation and Test in Europe, IEEE, 2009, pp. 682 – 687.

[LFP+10] P. Lokuciejewski, H. Falk, S. Plazar, P. Marwedel, and L. Thiele: *Multi-Objective Exploration of Compiler Optimizations for Real-Time Systems,* submitted to ASPLOS 2010

[PIEP09] P. Pop, V. Izosimov, P. Eles, and Z. Peng. Design Optimization of Time- and Cost-Constrained Fault-Tolerant Embedded Systems with Checkpointing and Replication. IEEE Transactions on Very Large Scale Integrated (VLSI) Systems, 17(3):389-402. 2009.

[WGRSPF09] R. Wilhelm, D. Grund, J. Reineke, M. Schlickling, M. Pister and C.Ferdinand: *Memory Hierarchies, Pipelines, and Buses for Future Architectures in Time-Critical Embedded Systems*, IEEE TCAD, July 2009

[WW09] S. Wilhelm and B. Wachter. *Symbolic state traversal for WCET analysis*. Proceedings of the International Conference on Embedded Software (EMSOFT), October 2009.

-- The above are new references, not present in the Y1 deliverale --

3.4 Keynotes, Workshops, Tutorials

Keynote: *Reinhard Wilhelm USAAR:* **Embedded Systems - Trends, Successes, Challenges**
10th Anniversary of the Hasso-Plattner Institute
Potsdam – Nov. 18, 2009
To celebrate its 10th anniversary, the Hasso-Plattner-Institiute holds a conference "Informatik-Impulse". Reinhard Wilhelm is invited to present an overview of the challenges in embedded systems world and to give an outlook onto future developments.

Invited talk: *Petru Eles, Linköping:* **Predictable Implementation of Real-Time Applications on Multiprocessor Systems on Chip***,*
Ninth International Workshop on Worst-Case Execution Time Analysis
Dublin, June 30, 2009.
With this occasion several results obtained in the ARTIST context have been made accessible to an international audience. They are related, in particular, to the predictability of applications running on multiprocessor systems with shared communication infrastructure.

Keynote : *Peter Puschner (TU Vienna):* **From Performance to Time-Predictability**
9th Architectures and Compilers for Embedded Systems (ACES) Symposium
Edegem, Belgium – September 7-8, 2009
This keynote outlined the problems of building predictable hardware/software systems and discussed strategies for constructing systems that provide both temporal predictability and performance.
http://www.elis.ugent.be/aces/index.php?page=activities

Keynote: *Reinhard Wilhelm USAAR:* **Timing Analysis and Timing Predictability**
Tag der Informatik
RWTH Aachen – December 4, 2009
Reinhard Wilhelm is invited to give a talk introducing timing analysis and timing predictability in embedded systems. The current challenges and existing timing analysis algorithms will be discussed as well as the additional challenges posed by multi-core systems and approaches to achieve predictability for them.
http://www.nets.rwth-aachen.de/content/current_events/tdi/pro/index.html

Keynote : *Reinhard Wilhelm USAAR:* **The PROMPT Design Principles for Predictable Multi-Core Architectures**
Software & Compilers for Embedded Systems (SCOPES) 2009
Nice, France – April 24, 2009
The presentation proposes design principles for multi-core architectures to provide efficiently predictable good worst-case performance as needed for embedded control in the aeronautics and automotive industries supporting the Integrated Modular Avionics (IMA) and the

Automotive Open System Architecture (AUTOSAR) development trends. One background challenge for the talks is that proving the correctness of a modern high-performance processor is beyond the reach of verification methods, and that even the chances to derive reliable and precise bounds on execution times are endangered by exactly these developments.
This talk presents a development process oriented at achieving predictability at all levels of the architecture hierarchy.
http://www.scopesconf.org/scopes-09

Keynote : *Lothar Thiele, ETHZ:* **Distributed Embedded Systems - Reconciling Computation, Communication and Resource Interaction**
15th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA 2009)
Peijing, China – August, 2009

Developing associated system design methods that lead to timing-predictable and efficient embedded systems imposes a major challenge. The presentation covered various attempts to address the above challenges, their use in embedded system design and various open issues.
http://www.cs.cityu.edu.hk/rtcsa2009/keynoteII_abs.htm

Invited Talk: *Reinhard Wilhelm USAAR:* **Predictable Multi-Cores**
Verimag – February 13, 2009.
Reinhard Wilhelm was invited to give this talk at a local colloquium at Verimag to explain the PROMPT design principles on predictable multi-core architectures.

Tutorial : *Peter Puschner (TU Vienna* **WCET Analysis: Problems, Methods and Time-Predictable Architectures**
Acaces 2009. Fifth International Summer School on Advanced Computer Architecture and Compilation for Embedded Systems
Terrassa (near Barcelona), Spain – July 12-18, 2009

The annual ACACES summer school is organized by the European Network of Excellence on High Performance and Embedded Architecture and Compilation (HiPEAC). The invited tutorial on WCET analysis and time-predictable hardware/software architectures was given by Peter Puschner (TU Vienna). A complete lecture of the tutorial was dedicated to techniques for constructing time-predictable real-time embedded systems.
http://www.hipeac.net/acaces2009/

Workshop : Reconciliating Performance and Predictability
ESWEEK
Grenoble, France – October, 2009

The RePP workshop was concerned with embedded systems that are characterized by efficiency requirements on the one hand and critical constraints on the other. Such systems occur in many industry-relevant embedded application domains such as avionics, automotive, railway systems, power plants, construction machinery, and robotics. The slides from the workshop provide interesting contributions to the area of characterization and realization of predictability. To this report, we also append the minutes from the discussion at the end of the workshop.

http://www.tik.ee.ethz.ch/~jchen/RePP/

Workshop : Software & Compilers for Embedded Systems (SCOPES) 2009
Nice, France – April 23-24, 2009

SCOPES focuses on the software generation process for modern embedded systems. Topics of interest include all aspects of the compilation process, starting with suitable modeling and specification techniques and programming languages for embedded systems. The emphasis of the workshop lies on code generation techniques for embedded processors. The exploitation of

specialized instruction set characteristics is as important as the development of new optimizations for embedded application domains. Cost criteria for the entire code generation and optimization process include runtime, timing predictability, energy dissipation, code size and others. Since today's embedded devices frequently consist of a multi-processor system-on-chip, the scope of this workshop is not limited to single-processor systems but particularly covers compilation techniques for MPSoC architectures.

SCOPES 2009 was organized by Heiko Falk, one of the essential contributors to ArtistDesign from TU Dortmund.

http://www.scopesconf.org/scopes-09

Workshop: Shape Analysis, Timing Analysis
PUMA Workshop
San Servolo Island,Venice – Oct. 5 and 6, 2009
Reinhard Wilhelm gave talks on Shape Analysis and of Timing Analysis at the PUMA workshop. This workshop is an annual event held by the PUMA graduate school of LMU and TU Munich.
< http://puma.in.tum.de/wiki/Venice_2009>

ARTIST Summer School in Europe 2009
Autrans – September 7-11, 2009
The Artist Summer School included invited talks by several researchers, including Luca Benini, Jan Beutel, Jan Reineke, Lothar Thiele and Reinhard Wilhelm.
<http://www.artist-embedded.org/artist/Overview,1633.html>

--- The above are new references, not present in the Y1 deliverale --

4. Overall Assessment and Vision for the Transversal Activity

4.1 Assessment for Year 2

The collaborations in the context of this activity have progressed successfully. The PREDATOR project, which involves several partners of this activity, has entered its second year and continued the successful work of Year 1. Examples of collaborations include the worst-case execution time aware compiler wcc (Dortmund, AbsInt, USAAR) and cache-aware scheduling (USAAR, SSSA). Other collaborations on hardware design for multi-core systems are underway, and should be strengthened in the next year.

The activity has achieved good results advancing on the way towards timing predictability on multi-core platforms. The detrimental effects of interferences on worst-case execution time and on predictability have been discussed, and we have provided means to quantify or reduce their effects in areas such as cache-related preemption.

On the problems of designing multi-core hardware that support predictability, significant progress is under way, especially concerning techniques to manage shared resources, most notably the shared buses. Several partners are addressing the design of strategies to schedule the bus for predictability, while also making more efficient use of chip resources. We also see that the design ideas for predictable multi-core hardware are slowly picked up by the industry. This shows that the projects have been successful in raising the awareness on predictability issues in the community.

Concerning the goal of reconciling timing analysis with compilation as stated in the second bullet of Section 1.3, the WCET-aware compiler WCC developed at TU Dortmund has matured in the past 12 months. The compiler is equipped with test-benches consisting of approx. 4000 benchmarks used to check the correctness of the compiler and to verify the effect of its WCET-aware optimizations. In ArtistDesign year 2, WCC has moved from an academic research software to a compiler ready for commercialization. In the next 12 months, it is foreseen to extend WCC towards code generation and optimization for multi-process systems, including real-time operating systems. For this purpose, collaboration between TU Dortmund, AbsInt and the ArtistDesign Cluster on Operating Systems and Networks will be strengthened

Our results on the PRET–C programming language and the dedicated ARPRET processor show that it is possible to reconcile time-predictability with excellent real-time performance. The evaluations we have performed with benchmarks written in PRET-C, Esterel, SyncCharts, and ProtoThreads show that PRET-C systematically outperforms the alternative approaches. We believe that this has been made possible thanks to our use of a reactive processor core inside the ARPRET architecture. Indeed, the instruction set of a reactive processor is dedicated to execute very efficiently reactive programs such as Esterel, with direct support for preemption, suspension, synchronization between threads and with the environment, and so on. This feature, which distinguishes our approach from concurrent ones (like that of UC Berkeley and Comlubia University), is the key to obtain both time-predictability and excellent real-time performances.

We can see a general interest in issues concerning the fundamental concepts relating to predictatiliby, such as characterizing and making more precise the concepts of predictability and robustness in different contexts. Another topic is that of designing more predictable, and even deterministic, system and computer architectures, which do not unnecessarily sacrifice performance features of modern computer architecture.

In the second year, we (co-)organized a workshop on predictability (RePP). The presentations and discussions at the workshop are documented, and will certainly inspire new research directions on this topic.

4.2 Overall Assesment since the start of the ArtistDesign NoE

We have performed solid work on identifying threads to predictability, quantifying their influences and giving guidances on designing more predictable systems. This work is most advanced for single-core systems, but with recent advances we also begin to handle the problems posed by multi-core systems. We are happy that work on the challenges of multi-core architectures seems to have gained speed, and interesting results should emerge during the course of ArtistDesign. As an example, work at TU Vienna indicates that time-sliced bus arbitration yielded predictable timing but not the desired performance. Further investigations into the memory systems are planned to get a better performance.

There are several collaborations between the partners under way. Not only is theoretical knowledge being spread across the partners, but the work is implemented into actual tools and tools of multiple partners are integrated, such as aiT, WCC, and MPARM.

-- The above is new text, not present in the Y1 deliverable --

4.3 Indicators for Integration

There were no explicit indicators stated for this activity for the second year in the Description of Work, but we can provide some observations below.

- 10 joint publications (a slight increase from Y1)
- A large number of research collaborations and visits involving more than one partner, including several European projects
- Successfully conducting a workshop on “reconciling predictability and performance”, which will inspire research in ArtistDesign and elsewhere on predictability.
- Educational events (summer schools) organized or co-organized by the partners. (Two ARTIST summer schools, and a strong presence at the summer school in Florianopolis

-- Changes wrt Y1 deliverable --

This is updated wrp. to year 1.

4.4 Long-Term Vision

Further work needs to be done to derive sound design principles for performance and predictability in system design. A key aspect in improving the predictability of embedded systems is to take a more holistic approach, considering all the levels of system design.

An overall goal should be to arrive at the definition of a multi-core architecture, including low-level hardware specification and scheduling, which can be subjected to timing analysis with a high degree of precision while still giving all the benefits of hardware parallelism. Within this

work, all areas of system architecture, operating system design, compilation and timing analysis should be integrated.

At the programmer level, compiler technology and programming language constructs should be developed that allow the developer to assess and take predictability into account directly in the software development. WCET-aware compilation, predictable use ot system features, and predictable programming constructs should be developed.

Similar development should be aimed for operating systems, on which applications can be run so that the timing of the whole software system is well-predictable. To this end the software structures of both, the applications taks and the operating systems will have to be predictable, and unwanted timing interactions (OS-task and task-task interactions) have to be avoided, or at least minimized.

The activities within ArtistDesign and connected activities should provide the conceptual tools for these developments, as well as provide reference implementations to push their realization in industry Having a prototypical implementation of, e.g., a timing aware compiler, a predictable multicore architecture, and an operating system would be considered a successful result.

-- Changes wrt Y1 deliverable --

Has been thoroughly revised in comparison with Y1.

5. Minutes from the RePP Workshop

Since the RePP workshop was co-organized by this activity, we include the minutes of the concluding discussion as part ot this activity report. They are here:

Reconciling Performance with Predictability

Claire Burguière Andreas Schranzhofer Lothar Thiele
Reinhard Wilhelm

1 Introduction

The RePP workshop is concerned with embedded systems that are characterized by e□ciency requirements on the one hand and critical constraints on the other. Such systems occur in many industry-relevant embedded application domains such as avionics, automotive, railway systems, power plants, construction machinery, and robotics.

O□ine guarantees for the satisfaction of critical constraints have to be derived by appropriate methods. The di□culty of deriving guarantees strongly depends on the predictability properties of systems, in particular of the employed processor architecture, the software design discipline, the operating system including the scheduling strategy, the communication mechanism, and the used middleware. However, at the same time, system e□ciency is measured by means of average-case behavior under di□erent criteria such as performance, utilization of resources, and power consumption.

Unfortunately, it can be observed that in computer system design the gap between average-case and worst-case behavior increases rapidly. The technical reasons for the limited time-predictability are well known, for example the variation and non-determinism of the system environment and the interference caused by the use of shared resources.

The aim of the workshop was to discuss approaches that attack the combination of the two goals, the improvements of worst-case predictability and of average-case performance, on all system layers and in the layering principle itself.

Lothar Thiele and Andreas Schranzhofer are with Computer Engineering and Networks Laboratory, ETH Zurich, Switzerland, {thiele,schranzhofer}@tik.ee.ethz.ch

†Reinhard Wilhelm and Claire Burgui`ere are with the Compiler Design Lab, Saarland University, Germany, {wilhelm,burguiere}@cs.uni-saarland.de

In the sections below we provide a summary of the main discussion of this workshop. The three main topics were:

- What is Predictability?
- How to reconcile predictability and performance?
- How to design predictable multi-core architecture?

2 Predictability

We are interested in a certain aspect of a software/hardware system's behavior, e.g. the time it takes to execute the software on the hardware, the energy it consumes, the maximally space it occupies etc.

Measurement: We may observe one execution of the software on the hardware by measuring this one aspect.

Determinism (`a la PRET): This property describes that an execution of the software always leads to the same value of this aspect.

In case of variability,

Analysis: We are more interested in statically analyzing what the variance across all executions of the software under this one aspect is.

Predictability: We are even more interested in determining how precise we can e□ciently predict this aspect of all executions of all programs on the given hardware platform.

Robustness: This property describes how big the change of this aspect is given a change in the system's state.

The definition of timing predictability and how it can be measured is closely related to the definition of determinism and robustness. System that do not comply to these properties should not be used for hard real time systems or critical systems. With these assumptions, bounds on the influence of a change in the system's state can be computed and therefore measures that represent the execution time variability, as proposed by Daniel Grund, are meaningful. This measure does not make any statement about the distribution of the actual execution traces within this variability window. In other words, for a particular system, the WCET might be very high and the BCET very low, resulting in a low predictability measure, while two other traces concentrate in a significantly smaller variability window. Conclusively, a single trace dominates the predictability measure. It is unclear how to take into account the distribution into the predictability measure and how this a□ects the expressiveness of the measure.

Another point of view, on how predictability can be defined, is to specify the hardness of the WCET analysis. The hardness to derive the WCET determines the predictability, i.e., systems that are easily analyzable are more predictable than systems that are hard to analyze. This definition overcomes one weakness of the previously discussed variability related measure. Namely, that depending on the tightness of the WCET analysis, the measure can become arbitrarily bad.

This targets predictability with respect to timing. However, predictability can also relate to value, energy consumption and many more. Predictability with respect to timing has already been achieved: 30 years ago, when computing platforms were simpler and mechanisms such as pipelining were not used.

As a summary, the following questions remain unanswered, or their respective answers remain disputed:

- Is it an option to make a step backwards and drop features such as pipelining, caches, shared resources etc., in order to achieve the next leap in hard real-time computing?
- Has predictability to be defined as a measure related to the inherent variability induced by the hardware platform?
- Is WCET the right measure for classifying real time systems at all?

3 Performance

As shown in the previous section, the discussion did not lead to a single definition of predictability. Similarly, the notion of performance was used with di□erent meanings. For instance, one can think in terms of average-case performance or worst-case performance.

Consider average-case performance, someone suggested that the variability of the execution time could have an influence on both predictability and performance. The more performant a processor, the more complex is its architecture, the more variability might be induced by this architecture and the less predictable is the execution time. That led to the main question of the discussion: Performance. How much performance do we want to trade for predictability? The answer seems to depend on the industrial partner involved (avionics, automotive,...). While predictability is of crucial importance for avionic systems, the willingness to trade performance for predictability might be much less for automotive systems.

Regarding worst-case performance, the following question has been raised: If systems are designed for worst-case performance, is there a loss of worst-case performance when going to a more predictable system?

Furthermore, in hard real-time systems the performance (average and worst-case) might be dominated by the memory performance. Then, in order to reconcile predictability and performance, it could be a starting point to use a predictable memory architecture.

Reconciling predictability and performance appeared to be a main topic in the design of multi-core architecture.

4 Multi-Core

Multi-Core and Multiprocessor Platforms, using shared resources such as memory hierarchies (cache, scratchpad memory, buses) seem to be the common ground to implement hard real time systems. E□ciency can be achieved by concurrent execution of independent tasks/applications, while access to shared resource causes additional delays. With an increasing number of cores in these systems, the question arises whether pipelines, caches etc. are still determinant factors in the variability of the WCET.

Some presentations featured a similar multi-core architecture. Each core has a private

first level (L1) cache. The next level (second level of cache or main memory) is shared and accessed through a shared bus or a cross-bar. Nevertheless, this was just one proposal. The right architecture and memory hierarchy is still an open question.

Many works focus on systems, where accesses to the shared resources are separated from general purpose computations. Task/Applications are separated in three phases, where in the first phase required data is acquired, in the second phase computations are performed, and in the third phase altered and/or new data is written to the shared resource. This model of accessing the shared resources allows to derive tighter bounds on the WCET, since the variability is restricted. This model is used in conjunction with time triggered architectures, that allow e□cient analysis, but require substantial hardware support.

Topics of special interest are the compositionality of the WCET analysis. In other words, composing a system by two analyzable systems should guarantee the resulting systems analyzability. How does the variability of the execution time propagate? Atomic units of computations might have a high variability between BCET and WCET, but the composition of a large quantity of them to a larger task might have a very low variability.

- What is the right architecture for predictable systems?
- Is resource sharing the dominant factor in the variability of the WCET in Multi-Core systems?
- What is the right memory and communication architecture in Multi-Core systems?
- How can timing analysis be made modular?

6. Transversal Activity Participants

-- Changes in the Cluster Participants wrt Y1 deliverable --

University of Kiel added as affiliated partner.

6.1 Core Partners

Team Leader **Leader for transversal activity "Design for predictability and Performance"**	
	Bengt Jonsson http://user.it.uu.se/~bengt/
Technical role(s) within ARTIST2	Participant in discussions, contributions regarding compositionality, modelling, analysis of timing properties, tool building (TIMES)
Research interests	Research interests include: embedded systems, semantics, verification, modelling, specification, testing of distributed and embedded systems
Role in leading conferences/journals/etc in the area	Have been PC member of most conferences in the area.
Notable past projects	ASTEC, Competence Center for Software Technology, 1995-2005. http://www.astec.uu.se/ WOODDES (IST project) A UML profile for Automotive industry http://wooddes.intranet.gr/ Advance http://www.liafa.jussieu.fr/~haberm/ADVANCE/ Regular model checking (www.regularmodelchecking.com)

Team Leader

	Prof. Luca Benini, University of Bologna http://www-micrel.deis.unibo.it/%7Ebenini/
Technical role(s) within ArtistDesign	Member of the Strategic Management Board Co-leads Hardware Platforms and MPSoC Design Participates in Intercluster activity: Design for Adaptivity Participates in Intercluster activity: Design for Predictability and Performance Leader of the JPRA Activity: “Platform and MPSoC Design”
Research interests	(i) Development of power modeling and estimation framework for systems-on-chip. (ii) Development of optimal allocation and scheduling techniques for energy-efficient mapping of multi-task applications onto multi-processor systems-on-chips. (iii) Development of energy-scavenging techniques for ultra-low power sensor network platforms.
Role in leading conferences/journals/etc in the area	▪ Program chair and vice-chair of Design Automation and Test in Europe Conference. ▪ Member of the 2003 MEDEA+ EDA roadmap committee 2003. ▪ Member of the IST Embedded System Technology Platform Initiative (ARTEMIS): working group on Design Methodologies ▪ Member of the Strategic Management Board of the ARTIST2 Network of excellence on Embedded Systems ▪ Member of the Advisory group on Computing Systems of the IST Embedded Systems Unit. ▪ Member of the technical program committee and organizing committee of several technical conferences, including the Design Automation Conference, International Symposium on Low Power Design, the Symposium on Hardware-Software Codesign. He is Associate Editor of the IEEE Transactions on Computer-Aided Design of Circuits and Systems and of the ACM Journal on Emerging Technologies in Computing Systems. ▪ Fellow of the IEEE.
Notable past projects	ICT-Project **REALITY - *Reliable and variability tolerant system-on-a-chip design in more-moore technologies.*** Funded under 7th FWP (Seventh Framework Programme). ICT-2007.3.1 Next-Generation Nanoelectronics Components and Electronics Integration. Start date: 01/01/2008; Duration: 30 months; Contract Type: Collaborative project; Project Reference: 216537; Project Cost: 4.45 million euro; Project Funding: 2.9 million euro. ICT-Project **PREDATOR - *Design for predictability and efficiency.*** Funded under 7th FWP (Seventh Framework Programme). ICT-2007.3.3 Embedded Systems Design. Start date: 01/02/2008; Duration: 36 months; Contract Type: Collaborative project; Project Reference: 216008; Project Cost: 3.93 million euro; Project Funding: 2.8 million euro.

	ICT-Project **GALAXY - *interface for complex digital system integration.*** Funded under 7th FWP (Seventh Framework Programme). ICT-2007.3.3 Embedded Systems Design. Start date: 01/12/2007; Duration: 36 months; Contract Type: Collaborative project; Project Reference: 214364; Project Cost: 4.08 million euro; Project Funding: 2.9 million euro. ICT-Project **DINAMICS - *Diagnostic Nanotech and Microtech Sensors.*** Funded under 6th FWP (Sixth Framework Programme). FP6-NMP 'Nanotechnologies and nanosciences, knowledge-based multifunctional materials and new production processes and devices'. Contract Type: Integrated project; Project Reference:IP 026804-2. Start date: 01/04/2007. Duration: 18 + 30 months. Project Cost:7276856 Euro. Project Funding: 4499542 Euro. http://www.dinamics-project.eu/ ICT-Project **SHARE - *Sharing open source software middleware to improve industry competitiveness in the embedded systems domain.*** Funded under 7th FWP (Seventh Framework Programme). ICT-2007.3.7 Network embedded and control systems. Start date: 01/05/2008; Duration: 24 months; Contract Type: Coordination and support actions; Project Reference: 224170; Project Cost: 1.1 million euro; Project Funding: 590000.00 euro.

Team Leader	
	Professor Alan Burns University of York, UK URL: www.cs.york.ac.uk/~burns
Technical role(s) within ArtistDesign	Undertakes research in real-time systems scheduling, particularly for flexible systems. Also concerned with the development of programming languages for this domain.
Research interests	Scheduling, languages, modeling and formal logics.
Role in leading conferences/journals/etc in the area	Previous Chair of the IEEE Technical Committee on Real-Time Systems. Edited special issue of ACM Transactions on Embedded Systems (on education).
Notable past projects	DIRC – Dependability Interdisciplinary Research Collaborations – A large, UK, 6-year, multisite project looking at dependability of computer-based systems. Burns was a PI and managed the work on temporal aspects of dependability. FIRST – EU funded project concerning flexible scheduling FRESCOR – EU follow on project to FIRST

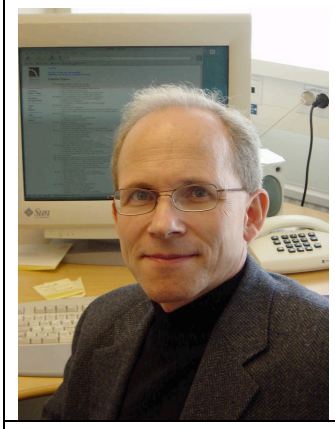	Petru Eles (Linköping University)
Technical role(s) within ArtistDesign	Main areas of research: Embedded Systems ArtistDesign activities and role: Communication centric systems, system analysis, optimisation, low power embedded systems, power management, modelling, analysis, and simulation of distributed embedded systems, predictable real-time systems, fault tolerance.
Research interests	Research interests include real-time systems, design of embedded systems, electronic design automation, hardware/software co-design,.
Role in leading conferences/journals/etc in the area	- Associate Editor, IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems; - Associate Editor, IEE Proceedings - Computers and Digital Techniques; - TPC Chair and General Chair, IEEE/ACM/IFIP International Conference on Hardware/Software Codesign and System Synthesis (CODES/ISSS). - Topic chair, Design Automation and Test in Europe (DATE). - Topic Chair, Int. Conference on Computer Aided Design (ICCAD). - Program chair of the Hw/Sw Codesign track, IEEE Real-Time Systems Symposium (RTSS). - TPC Chair IEEE Workshop on Embedded Systems for Real-Time Multimedia (ESTIMedia). - Steering Committee Chair, IEEE/ACM/IFIP International Conference on Hardware/Software Codesign and System Synthesis (CODES/ISSS).
Awards / Decorations	- Best paper award, European Design Automation Conference (EURO-DAC), 1992. - Best paper award, European Design Automation Conference (EURO-DAC), 1994. - Best paper award, Design Automation and Test in Europe (DATE), 2005. - Best presentation award, IEEE/ACM/IFIP International Conference on Hardware/Software Codesign and System Synthesis

	(CODES/ISSS), 2003. - IEEE Circuits and Systems Society Distinguished Lecturer, for 2004 - 2005.

Partner in Activitiy on Predictability	
	Alain Girault (INRIA Grenoble Rhône-Alpes)
Technical role(s) within ArtistDesign	Main areas of research: Embedded Systems ArtistDesign activities and role: formal methods for the design of embedded systems, predictable real-time systems, dependability analysis and design, fault tolerance.
Research interests	Research interests include embedded and real-time systems, formal methods, dependability, fault tolerance.
Role in leading conferences/journals/etc in the area	- Associate Editor, Eurasip Journal on Embedded Systems; - TCP co-chair of the Workshop on Model-driven High-level Programming of Embedded Systems (SLA++P'08).

Partner in Activitiy on Predictability	
	Michael González Harbour (Universidad de Cantabria)
Technical role(s) within ARTISTDesign	ArtistDesign activities and role: Participates in Operating Systems and Networks cluster, and also in the Intercluster activity: Design for Predictability and Performance
Research interests	Research interests include schedulability analysis for distributed

	real-time systems, real-time operating systems, real-time languages
Role in leading conferences/journals/etc in the area	Has been program committee chair in the ECRTS and Ada-Europe conferences, and in the International Real-Time Ada Workshop. Has participated in the past five years in the program committes of the following international conferences: ECRTS, Ada-Europe, RTSS, RTAS, ACM Symposium on Applied Computing, WPDRTS, CORDIE, DATE, ETFA, EUC, EMSOFT, IRTAW, EDF. Has been invited editor in the Real-Time Systems Journal and the Eurasip Journal on Embedded Systems. Has participated actively in the development of the POSIX standards, in the extensions of operating systems services for real-time applications.
Notable past projects	FRESCOR: Framework for Real-time Embedded Systems based on COntRacts (EU project) FIRST: Flexible Integrated Real-Time Systems Technology (EU project)
Further Information	Group home page: http://www.ctr.unican.es MAST toolset: http://mast.unican.es MaRTE OS: http://marte.unican.es

Cluster Leader **Activity Leader for "Software Synthesis and Code Generation"**	
	Prof. Dr. Peter Marwedel (TU Dortmund) http://ls12-www.cs.tu-dortmund.de/~marwedel/
Technical role(s) within ArtistDesign	Cluster leader, activity leader SW Synthesis and Code Generation Improved code quality for embedded applications is the main goal of the work at Dortmund University. Due to the widening gap between processor and memory speeds, emphasis has been on improving the efficiency of memory accesses, in terms of average and worst case execution time and in terms of the energy consumption.
Research interests	Peter Marwedel's Embedded Systems Group focuses on embedded software. Particular emphasis is on compilers for embedded processors. One of the very first publications in this area, the book "Compilers for Embedded Processors", edited by Peter Marwedel and Gert Goossens, was the result of the CHIPS project, funded by the European Commission. The group's current focus is on advanced optimizations for embedded processors (e.g. by using bit-level data flow analysis) and energy-aware compilation techniques.

	Current research also includes high-level transformations of algorithms as well as WCET-aware code generation.
Role in leading conferences/journals/etc in the area	Member of the EDAA (European Design and Automation Association) Main Board. Editorial Board Member of the Journal of Embedded Computing. Editorial Board Member of the Microelectronics Journal. Co-Founder and Steering Board Chair of the SCOPES Workshop (Software and Compilers for Embedded Systems) Series. >14 years of service for the DATE conference and its predecessors (program chair: 3 times, chairman of the steering committee, European representative to ASPDAC) DAC: Topic chair and reviewer Various other conferences
Notable past projects	MAMS: Multi-Access modular-services framework, national project funded by the German Federal Ministry of Education and Research (BMBF) MORE: Network-centric Middleware for group communications and resource sharing across heterogeneous embedded systems, supported by the European Commission http://www.ist-more.org HiPEAC: European NoE on High-Performance Embedded Architecture and Compilation; http://www.hipeac.net Others: Various earlier projects supported by the EC, DFG etc.
Awards / Decorations	Teaching award, TU Dortmund, 2003 DATE fellow, 2008
Further Information	CEO of the Informatik Centrum Dortmund (ICD), a technology transfer centre founded in 1989.

	Dr. Stylianos Mamagkakis (IMEC vzw.) http://www.imec.be
Technical role(s) within	SW Synthesis and Code Generation; collaboration with TU

ArtistDesign	Dortmund on high-level transformations for source code optimizations.
Research interests	Stylianos Mamagkakis received his Master and Ph.D. degree in Electrical and Computer Engineering from the Democritus Uni. Thrace (Greece) in 2004 and 2007, respectively. Since 2006, he coordinates a team of PhD students within the NES division at IMEC, Leuven, Belgium. His research activities mainly belong to the field of system-level exploration, with emphasis on dynamic resource management and system integration.
Role in leading conferences/journals/etc in the area	Stylianos Mamagkakis has published more than 25 papers in International Journals and Conferences. He was investigator in 6 research projects in the embedded systems domain funded from the EC as well as national governments and industry.
Notable past projects	EASY IST project Energy-Aware System-on-chip design of the HIPERLAN/2 standard, http://easy.intranet.gr/ AMDREL IST project Architectures and Methodologies for Dynamic Reconfigurable Logic, http://vlsi.ee.duth.gr/amdrel/

Transversal Activity Leader **Activity Leader for "NoE Integration: Low Power"**	
	Prof. Dr. Peter Puschner (TU Vienna) Real-Time Systems Group Institute of Computer Engineering Vienna University of Technology http://www.vmars.tuwien.ac.at/people/puschner.html
Technical role(s) within ArtistDesign	Peter Puschner and his group are participating in the timing analysis and design for predictability activities of ArtistDesign. They will provide technical contributions in compiler support for timing analysis, software/hardware architectures that make real-time systems more time-predictable and composable, and operating systems with predictable timing.
Research interests	Peter Puschner's main research interst is on real-time systems. Within this area he focuses on Worst-Case Execution Time Analysis and Time-Predictable Architectures.

Role in leading conferences/journals/etc in the area	Member of the Euromicro Technical Committee on Real-Time Systems, the steering committee of the Euromicro Conference on Real-Time Systems (ECRTS) Member of the advisory board and organizers committee of the IEEE International Symposium on Object- and Component-Oriented Distributed Computing (ISORC) conference series Chair of the Steering Committee of the Euromicro Workshop on Worst-Case Execution-Time Analysis (WCET) series
Notable past projects	DECOS - Dependable Embedded Components and Systems Develop the basic enabling technology to move from a federated distributed architecture to an integrated distributed architecture. http://www.decos.at MoDECS - Model-Based Development of Distributed Embedded Control Systems Model-based construction of distributed embedded control systems: shift from a platform-oriented towards a domain-oriented, platform-independent development of composable, distributed embedded control systems. http://www.modecs.cc NEXT TTA Enhance the structure, functionality and dependability of the time-triggered architecture (TTA) to meet the cost structure of the automotive industry, while satisfying the rigorous safety requirements of the aerospace industry. http://www.vmars.tuwien.ac.at/projects/nexttta/
Awards / Decorations	
Further Information	

Team Leader	
	Alberto Sangiovanni Vincentelli (PARADES) http://www.parades.rm.cnr.it
Technical role(s) within ARTIST2	Bring in Expertise in embedded system modelling, validation, tools and methodologies and IC design.

	Deep involvement in cooperation with the industry: tools (co-founder Cadence and Synopsys), telecommunications (Telecom Italia), automotive (member of the GM STAB)
Research interests	Embedded system design methodologies and tools including modelling, validation, synthesis and formal verification, semantic foundations.
Role in leading conferences/journals/etc in the area	Program Committee Member CODES and EMSOFT. Member of the Editorial Boards Member of the ARTEMIS High-level Group and Steering Committee
Notable projects	SPEEDS - Speculative and Exploratory Design in Systems Engineering Provide a semantics based modelling methods with analysing techniques to support the construction of complex embedded systems by composing heterogeneous subsystems together with a speculative tool-supported design process. HYCON NoE: Taming Hybrid Systems Center for Hybrid and Embedded Software Systems (CHESS) co-director Gigascale System Research Center, Core theme leader RIMACS: Industrial Automation
Awards/Decorations	IEEE Fellow, Member National Academy of Engineering, Kaufmann Award for pioneering contributions to EDA, IEEE Graduate Teaching Award, Gulliemin-Cauer Award, Darlington Award, Aristotle Award, University of California Distinguished Teaching Award

Participant in Activitiy on Adaptivity

	Lothar Thiele (ETH Zurich)
Technical role(s) within ARTISTDesign	Main areas of research: Embedded Systems and Software Artist2 activities and role: Communication Centric Systems: Formal Performance Analysis, Linking Simulation and Verification, Design Space Exploration of Embedded Systems
Research interests	Research interests include models, methods and software tools for the design of embedded systems, embedded software and bioinspired optimization techniques.
Awards / Decorations	In 1986 he received the "Dissertation Award" of the Technical University of Munich, in 1987, the "Outstanding Young Author Award" of the IEEE Circuits and Systems Society, in 1988, the Browder J. Thompson Memorial Award of the IEEE, and in 2000-2001, the "IBM Faculty Partnership Award". In 2004, he joined the German Academy of Natural Scientists Leopoldina. In 2005-2006, he was the recipient of the Honorary Blaise Pascal Chair of University Leiden, The Netherlands. Chair of ACM SIGBED.

	Prof. Dr. Dr. h. c. mult Reinhard Wilhelm (Saarland University) http://rw4.cs.uni-sb.de/people/wilhelm
Technical role(s) within ArtistDesign	Timing Analysis
Research interests	Compilers, Static Analysis, Timing Analysis
Role in leading conferences/journals/etc in the area	PC member of SCOPES, LCTES, MEMOCODE, RTSS etc. Steering committee member of EMSOFT, member at large of the steering committee of LCTES Member of the ACM SIGBED Executive Committee
Notable past projects	DAEDALUS
Awards / Decorations	Prix Gay-Lussac-Humboldt in 2007

	Honorary doctorates of RWTH Aachen and Tartu University in 2008 Konrad-Zuse Medal in 2009
Further Information	Co-founder of AbsInt Angewandte Informatik GmbH Scientific Director of the Leibniz Center for Informatics Schloss Dagstuhl

6.2 Affiliated Academic Partners

Team Leader	
	Rolf Ernst (TU Braunschweig)
Technical role(s) within ArtistDesign	Main areas of research: Embedded Systems Participates in Hardware Platforms and MPSoC Design Participates in Intercluster activity: Design for Adaptivity Participates in Intercluster activity: Design for Predictability and Performance Participates in Intercluster activity: Integration Driven by Industrial Applications
Research interests	Research interests include embedded architectures, hardware-/software co-design, real-time systems, and embedded systems engineering.
Role in leading conferences/journals/etc in the area	He chaired major international events, such as the International Conference on Computer Aided Design of VLSI (ICCAD), or the Design Automation and Test in Europe (DATE) Conference and Exhibition, and was Chair of the European Design Automation Association (EDAA), which is the main sponsor of DATE. He is a founding member of the ACM Special Interest Group on Embedded System Design (SIGBED), and was a member of the first board of directors. He is an elected member (Fachkollegiat) and Deputy Spokesperson of the "Computer Science" review board of the German DFG (corresponds to NSF). He is an advisor to the German Ministry of Economics and Technology for the high-tech entrepreneurship program EXIST (www.exist.org).

	Prof. Dr. Reinhard von Hanxleden (Kiel University) http://www.informatik.uni-kiel.de/en/rtsys/contact/
Technical role(s) within ArtistDesign	Affiliated to "Design for Predictability and Performance".
Research interests	• Model-based design of complex system, modeling pragmatics, automatic synthesis/layout of graphical models, Eclipse. • Synchronous languages, embedding reactive control flow into classical programming languages (C/C++/Java). • Reactive/predictable processor design.
Role in leading conferences/journals/etc in the area	• PC member of APGES, SLA++P, SRDS, TCMC. • Co-organizer of Synchronous Programming (SYNCHRON).
Notable past projects	Dependable Embedded Components and Systems (DECOS) • EU 6th Framework Integrated Project • http://www.decos.at Model-Based Engineering of Electronic Railway Control Centers (MENGES) • EU regional funding
Awards / Decorations	• Teaching awards, CS at Kiel University, 2008 (3rd) and 2009 (2nd)