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

Transversal Activity Progress Report for Year 3

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	February 4th 2011

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 Y2 deliverable --

No changes with respect to Year 2.

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 Y2 deliverable --

No changes with respect to Year 2.

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 Y2 deliverable --

No changes with respect to Year 2.

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 Y2 deliverable --

No changes with respect to Year 2.

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 Y2 deliverable --

No changes with respect to Year 2.

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 Y2 deliverable --

No changes with respect to Year 2.

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), Saarland University (Reinhard Wilhelm), and TU Dortmund (Peter Marwedel), to integrate worst-case execution time analysis, modular performance anlaysis and compiler development, in order to provide the programmer with direct information about predictability measures (in particular WCET) on the code that is being developed. One very interesting outcome of this collaboration is the predictability-aware compilers WCC, which is an execellent basis for driving and evaluating continued work on compiler and architectural techniques to achieve predictability. Important work can be performed 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 U Trento (previously PARADES) are collaborating on compositional languages and approach to embedded system that can be seen as an 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 scheduling techniques and bus access policies on multiprocessor platforms, with the goal to design predictable, yet efficient, designs. Partners include UoB, ETHZ, SSSA Pisa, Linköping, Braunschweig, DTU, TU Vienna, and Uppsala.

Partners are maintaining and ongoing discussion on basic concepts, aiming to provide a more firmly based definition of "predictability" and relating that to developed techniques. An important forum for such discussions has been the RePP Workshop, which was held during ESWEEK, Grenoble, October 2009, and which will again appear at DATE in Grenoble, March, 2011.

Finally, several partners (INRIA and U Kiel, and also with non ArtistDesign partners, U Auckland and UC Berkeley) are collaborating on time predictable execution platforms inspired by former reactive processors, which were tailored to the execution of the ESTEREL synchronous programming language. By taking advantage of the lock step execution provided by reactive processors, it will be possible to execute synchronous programs with both a high throughput and a precise WCET.

-- Changes wrt Y2 deliverable --

Updated to reflect development

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.

U Bologna, TU Dortmund, and TU 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 and the ACACES summer school).

TU Dortmund promotes education in embedded systems through a published textbook ("Embedded System Design"). Several other groups were asked to comment on the upcoming second edition of the book, to be published in 2011. This group also organizes the WESE workshop on embedded system education (a satellite workshop of ESWEEK). 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 workshop on the mapping of applications to MPSoCs and the workshop on software synthesis, held during ESWEEK in 2010.

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 Y2 deliverable --

Updated to reflect development

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 (see [Hen08] in Deliverable 15-(7.2)-Y1). The technical contribution was to suggest how predictability can be formalized as a form of determinism, and robustness as a form of continuity. IST Austria, together with VERIMAG, continued the effort on studying robustness and predictability in embedded systems in Year 3, following several directions:

- Hierarchical Timing Language (HTL) is a real-time coordination language for distributed control systems. We continued the work started in Year 2 and reported in Deliverables D6-(3.2)-Y2 and D15-(7.2)-Y2 and proposed a generalization of the virtual machine that executes HTL programs in order to support their hierarchical structure [GIKHV10]. This work that is a result of the collaboration with University of California at Berkeley and Uni. Trento, is presented in more details in D6-(3.2)-Y3.
- We extended our work on robust synthesis presented in D6-(3.1)-Y2 and D15-(7.2)-Y2 in order to support robust synthesis in presence of liveness [BCG+10]

Methodology for designing component-based real-time systems (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 fulfilment 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 OMG Standards (Cantabria)
Activities in support of two relevant standardization processes have been performed by UC at the OMG standardization body.

- In support of SysML, Cantabria has joined and participated in the discussions of the current sysML 1.3 Revision Task force.
- In support of the UML Profile for MARTE, in the context of the second revision task force, Cantabria has been responsible for the resolution of issues related to two

chapters of the MARTE 1.2 specification: Generic Resource Modeling and Schedulability Analysis Modelling. It has prepared and successfully introduced resolutions for all the pending issues in these two chapters as well as others related to them in the Generic Quantitative Analysis Modelling. This activity may be followed by looking at the wiki page of the task force: http://www.omgwiki.org/marte1.2. Also, with the support of the FP7 ADAMS support action, UC has built a new collaborative webpage for the standard using Drupan technology. Now, with the support of ArtistDesign, and in collaboration with CEA, Cantabria is in charge of its maintenance. This page is now the official OMG web page for the standard (http://www.omgmarge.org). In the last OMG meeting Cantabria has continue this effort by joining the new MARTE 1.3 Revision Task Force.

As in the last two years, this year Cantabria is also contributing as part of the Program Committee in the organization of the OMG's Workshop on Real-time, Embedded and Enterprise-Scale Time-Critical Systems.

Predictable, Fault-tolerant Embedded Systems Design: Fault Tolerant Applications With Hard and Soft Real-Time Constraints (Linköping, DTU)
In this work, we propose an approach for scheduling of predictable, fault-tolerant applications composed of soft and hard realtime processes running on distributed systems. We use process re-execution to tolerate transient faults. We propose a quasi-static scheduling algorithm that generates off-line a tree of fault-tolerant schedules that maximize the quality-ofservice of the application and, at the same time, guarantee deadlines for hard processes. At run time, the online scheduler, with very low online overhead, would select the appropriate schedule based on the occurrence of faults and the actual execution times of processes. The proposed approach can be useful on any distributed system whose worst-case communication delays can be obtained. Examples of such systems include multimedia systems in mobile phones, media players, TV-sets, and automotive infotainment. The approach is also applicable for safety-critical systems used, for example, in factory automation or automobiles. Prof. Paul Pop from DTU has visited Linköping with several occasions during this period.

Timing analysis and Compiler Techniques

Advances were made on the problem of time- predictable programs, by advancing the scope of static WCET analysis, and continuing the integration into the WCET-aware compiler.

Timing Analysis and Timing Predictability (USaar)
We have continued the work on cache replacement analysis: First, we devised a pair of precise and efficient must- and may-analyses for FIFO replacement, based on the novel principle of static phase detection. These analyses are more precise than the first ones for FIFO, which were reported on last year. Second, we addressed PLRU replacement, for which the precision of analyses still lags behind. We devised a more precise, but still expensive, must-analysis that can predict more cache hits than state-of-the art analyses.

We also continued the work on predictability. In one article we argue that some architectural features make timing analysis very hard, if not infeasible, but also shows how smart configuration of existing complex architectures can alleviate this problem. Furthermore, we resumed efforts towards a formal definition of predictability, arriving at definitions distinguishing between several kinds of predictability with respect to several sources of uncertainty.

We further elaborated on two techniques developed in 2009 to handle the additional uncertainty of cache effects arising from dynamic memory allocation. Tests with our predictable memory allocator confirm that in practice cache-awareness does not further increase worst-case execution times or fragmentation of real-time dynamic memory (de)allocation. Our

algorithm for transforming dynamic allocation into static allocation with comparable memory consumption was extended to cope with programs for which only parametrical bounds on the number of allocated objects are statically available.

WCET Analysis in the Presence of Context Switches (USaar, AbsInt, SSSA)
In preemptive real-time systems, scheduling analyses need—in addition to the worst-case execution time—the context-switch costs. The delay caused by cache misses due to context switches, which constitutes the major part of the context-switch costs, is referred to as *cache-related preemption delay* (CRPD). We developed a new approach to compute tight bounds on the CRPD for LRU set-associative caches, based on analyses of both the preempted and the preempting task. As the basis of our approach we introduce the notion of *resilience* of a memory block of an preempted task: the maximal number of memory accesses a preempting task could perform without causing an additional miss to this block. By computing lower bounds on the resilience of blocks and an upper bound on the number of accesses by a preempting task, one can guarantee that some blocks may not contribute to the CRPD. This resilience-based CRPD analysis considerably outperforms previous approaches. Ongoing work at Absint aims at integrating the resilience analysis in the timing analysis tool aiT.

Integration of timing analysis and compilation (TU Dortmund, AbsInt, ETHZ)
The work on Dortmund's Worst-Case Execution Time-aware C Compiler WCC has been continued in ArtistDesign Year 3.

TU Dortmund proposed a new WCET-driven cache-aware memory content selection algorithm, which allocates functions whose WCET highly benefits from a cached execution to cached memory areas. Vice versa, rarely used functions which do not benefit from a cached execution are allocated to non-cached memory areas [PLM10].

We exploited the concept of superblocks (regions in a program code that consist of multiple basic blocks) for the optimization of real-time systems. Superblock formation is now based on a novel trace selection algorithm which is driven by WCET data. Moreover, superblocks are translated for the first time from assembly to source code level. An adaption of the traditional optimizations common subexpression and dead code elimination to the proposed WCET-aware superblocks allows an effective WCET reduction [LKM10]. This activity was awarded three times, e.g. as best computer science thesis in Germany.

TU Dortmund developed techniques for WCET-aware basic block reordering in order to avoid unconditional branches and to support prediction of conditional branches. A genetic approach determines an optimized order of basic blocks of a function by applying evolutionary algorithms considering the WCET of the program to optimize as fitness value. Additionally, an integer-linear programming-based approach determines the optimal order of basic blocks and also takes the branch prediction into account.

A generic framework was developed by TU Dortmund, based on machine learning and evolutionary algorithms. Six different machine learning algorithms were systematically analyzed, and optimization of the different learner's parameter settings was done using evolutionary algorithms. In contrast to the majority of other approaches aiming at a reduction of the average-case execution time (ACET), the goal of this activity is the minimization of the worst-case execution time (WCET). Based on the proposed machine-learning framework, the well-known optimization loop-invariant code motion has been made WCET-aware [LSMM10].

Register allocation is considered the most important compiler optimization at all. TU Dortmund developed a novel WCET-aware register allocator based on integer-linear programming (ILP). The proposed ILP formulation relies on an already known approach for ILP-based register allocation minimizing spill code. Significant effort has been spent in order to model the WCET

and its heavily dynamic nature w.r.t. the worst-case execution path (WCEP) inside the ILP. Furthermore, care has been taken in order to model the impact of individual spill instructions generated by the register allocator on the processor's pipeline, as these additional instructions might contribute to the WCETs of their belonging basic blocks and thus might impact a program's WCEP and WCET.

Due to the complex interactions between a compiler's optimizations, the choice for a promising sequence of code transformations is not trivial. Compiler developers address this problem by proposing standard optimization levels, e.g., O3 or Os. However, previous studies have shown that these standard levels often miss optimization potential or might even result in performance degradation. In this activity, TU Dortmund and ETH Zurich propose the first adaptive WCET-aware compiler framework for an automatic search of compiler optimization sequences which yield highly optimized code. Besides the objective functions ACET and code size, we consider the WCET. To find suitable trade-offs between these objectives, stochastic evolutionary multi-objective algorithms identifying Pareto optimal solutions are exploited [LPFM+10a, LPFM+10b].

Since the WCET-analyzer integrated into the WCC is aimed at determining the WCET of isolated tasks, its application in scenarios with multiple tasks being scheduled on a single processor is limited. This is primarily due to the mutually invisible effects on the hardware states which largely affect each task's WCET. To perform safe estimations of the respective WCET of each task, particularly detailed knowledge of the cache states throughout the schedule is mandatory. Therefore, a framework for cache analysis for one or multiple tasks has been developed which is portable, modular, fast, precise and features state-of-the-art analyses for cache-related effects on the tasks' WCET. In the future, this framework will be the basis for scheduling- and operating-system-aware WCET optimizations integrated into WCC.

To support the ongoing efforts at TU Dortmund to feature WCET-aware multi-task optimizations, the ERIKA operating system provided by the University of Pisa (Giorgio Buttazzo) has been ported to the TriCore platform which plays an important role in today's automotive industry. To meet the industry's demand for highly time-predictable software, the WCC compiler not only supports automatic WCET-aware optimizations for single, isolated tasks but is also increasingly adopted to support these kinds of optimizations in the context of multiple tasks scheduled on a single processor. The TriCore-specific additions to ERIKA are officially packaged with an entirely open development platform provided by Evidence SRL.

In the past, WCET-aware scratchpad (SPM) allocation techniques have only been studied for single-task systems up to now. Currently, TU Dortmund extends the previous work on ILP-based SPM allocation for single-task systems towards multi-task systems. The totally available SPM capacity is partitioned into disjoint areas which can be used exclusively by the given tasks. A shared SPM region is available to all tasks and thus needs to be reloaded at context switch time. Based on the ILP formulations developed at TU Dortmund for single-task systems, the sizes of the individual SPM regions per task are computed, and the content of the tasks is distributed among these different regions.

Overall, the WCC compiler into which all above-mentioned techniques are integrated, can be considered the leading WCET-aware compiler. The integrated tool set allows studying the impact of optimizations for WCET minimization. WCC is currently evaluated at Bosch and is now able to generate and optimize industrial code representing an engine control system. First results indicate that WCC produces code with a WCET of a factor of 2 below that of the GCC compiler.

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

Statistical Analysis of Controller Area Network Message Response Times

Participants: UC Berkeley, GM, Trento

Modern automobile architectures are composed by tens of Electronics Control Units (ECUs) connected by several buses, most of which are Controller Area Networks (CAN). The availability of multiple ECUs can be exploited by distributing control tasks of one domain (for example, power train) to several ECUs. In this case, a number of distributed functions are assigned to multiple tasks executing concurrently on different modules and communicating via messages transmitted on CAN. Distributed functions include time-critical controls, but most often, also functions that are characterized by requirements for average performance together with hard deadline constraints (as for most active-safety functions) and functions with soft real-time requirements (controls for enhanced driver comfort). The definition of a new architecture framework for one or more car product families is an extremely important step: ECUs, networks and the topology of connections must be defined and frozen years in advance of production. Later, during the architecture lifespan, functions are placed on ECUs and communication scheduled on the bus. In [ZDGSV10] we presented a statistical approach to the early evaluation and selection of distributed embedded architectures for next-generation automotive controls, where the application performance depends on the end-to-end latencies of active-safety functions. Automobile architecture must be evaluated and selected having in mind that they will have a lifespan of 5 to 10 years and that during this lifespan the communication and computation load is partly unknown because new functions are still being decided on and have not been designed as yet. Hence, when verifying that the architecture is sufficiently robust with respect to constraints on latency and performance targets of present and future functionalities, loads can only be roughly estimated by looking at past trends or by exploiting early indications of designers. In this work, we considered an application model that is currently deployed in GeneralMotors E/E architectures and is supported by the AUTOSAR standard. We described the use of statistical analysis to compute the probability distribution of Controller Area Network (CAN) message response times when only partial information is available about the electrical architecture of a vehicle as well as about its functionality. We provided results that showed our statistical inference allows predicting accurately the distribution of the response time of a CAN message, once its priority has been assigned, from limited information such as the bus utilization of higher priority messages.

The results of the research obtained the best paper award at the IEEE Symposium on Industrial Embedded Systems and were selected for publication on the IEEE Transactions on Industrial Informatics.

OS/MW/Networks

Advances were made on achieving predictability in multiprocessor systems, concerning scheduling and management of shared resources.

Modelling and Analysis of Multiprocessor Systems with Shared Resources (TU Braunschweig) *(see also Activity on Platform and MPSoC Analysis)*
The classic assumption of many schedulability tests is that tasks' wost-case execution times are known, assumption that is not valid in case of multiprocessor systems where tasks mapped on different processors share common resources, e.g shared memories. In such setups, the worst-case execution times depend on the load imposed on the shared resource by other tasks in the system. TU Braunschweig has further worked on modelling and analysis methods for multiprocessor systems with shared resources. A method that captures more accurately the load imposed by tasks on the shared resources and therewith the interference between different cores that share common resources [SNE10] has been developed and applied to different multi-core setups to investigate the advantages on the obtained analysis results.

Predictable and efficient non-preemptive scheduling of multi-task applications (University of Bologna) *(see also Activity on Platform and MPSoC Analysis)*
Effective scheduling on multicore computing is a hard challenge as the goal is maximize the utilization of computational resources that can be integrated on a chip. Predictability is one of the most important issue for embedded multicore platforms. Many embedded applications run under real-time constraints, i.e. deadlines that have to be met for any possible execution. Predictability and computational efficiency are often conflicting objectives, as many performance enhancement techniques aim at boosting expected execution time, without considering potentially adverse consequences on worst-case execution. Hence applications with strong predictability requirements often tend to under-utilize hardware resources (e.g. forbidding or restricting the use of cache memories, limiting resource sharing, etc.). University of Bologna (UNIBO) developed a predictable and efficient non-preemptive scheduling of multi-task applications with inter-task dependencies. The focus is on a robust scheduling algorithm that proactively inserts additional inter-task dependencies only when required to meet the deadline for any possible combination of task execution times within the specified intervals. The approach does not need timers and related interrupts, since it avoids idle time insertion. UNIBO also proposes an iterative version of the algorithm for computing the tightest deadline that can be met in a robust way.

Enforcing Time-Predictability on Contemporary Computer Architectures (TU Vienna)
The dynamic decisions about resource allocation in contemporary ES computer architectures hcause 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 continued our work on strategies for eliminating timing anomalies by different code-generation techniques: Besides the insertion of nop instructions resp. instructions that create new register-use dependencies in the code thus removing anomalous timing effects we also studied techniques to eliminate dynamic decisions in branch predictors. The work on manipulating code to make the timing of branch predictors predictable could only be feasibly done for very simple two-bit predictors. For more complex predictors the increase in code size was simply too big. A paper on this work has been submitted for publication in a journal.

Predictable reliability and power consumption in distributed embedded systems (INRIA, U. Batna, and U. Casablanca)
We have developped a scheduling heuristics that, from a given software application graph and a given multiprocessor architecture, produces a static multiprocessor schedule that optimizes three criteria: its length (crucial for real-time systems), its reliability (crucial for dependable systems), and its power consumption (crucial for autonomous systems). Our tricriteria scheduling heuristics, TSH, uses the active replication of the operations and the data-dependencies to increase the reliability, and uses dynamic voltage scaling (DVS) to lower the power consumption. Thanks to a smart cost function, we generate schedules with a predictable reliability (at least that given as an objective by the user) and a predictable power consumption (at most that given as an objective by the user).

Energy intake prediction algorithms for systems powered by energy harvesters (University of Bologna)
Small size photovoltaic modules can harvest enough energy to power many personal devices and wireless sensor nodes. The prediction of solar energy intake is possible thanks to the periodical availability of the sunlight and its cyclic behavior. Thus, smart and innovative power management strategies can take advantage from intake prediction algorithms to optimize the energy usage by keeping the system in low power state as long as possible. On the other hand, very accurate predictions need time and energy because of complex calculations, thus

an algorithm that can provide the optimal trade-off between computational effort and accuracy is a breakthrough for systems with tight power constraints. University of Bologna introduced an innovative, efficient and reliable solar prediction algorithm based on the results achieved in **Y2**. The algorithm has been further enhanced to increase performance using a phase displacement regulator (PDR) which reduces the average error to less than 9.2% at a minimum energy cost. The proposed new algorithm compares favorably with several competing approaches.

Architecture and System Design: Multicore processors

Increasingly much work is devoted to achieving predictability on multicore architectures, including a large number of teams in the ARTISTDesign predictability activity.

Predictable Multi-core Communication (ETHZ with affiliated partner)
The work has been described in [SPCTC10], [PSCCT10] and [SCT10]. Multi-processor and multi-core systems are becoming increasingly important in time critical systems. Shared resources, such as shared memory or communication buses are used to share data and read sensors. We consider real-time tasks constituted by superblocks, which can be executed sequentially or by a time triggered static schedule.Three models to access shared resources are explored: (1) the dedicated access model, in which accesses happen only in dedicated phases, (2) the general access model, in which accesses could happen at anytime, and (3) the hybrid access model, combining the dedicated and general access model. For resource access based on a Time Division Multiple Access (TDMA) protocol, we analyze the worst-case completion time for a superblock, derive worst-case response times for tasks and obtain the relation of schedulability between different models. We conclude with proposing the dedicated sequential model as the model of choice for time critical resource sharing multi-processor/multi-core systems.

Employing COTS components in real-time embedded systems leads to timing challenges. When multiple CPU cores and DMA peripherals run simultaneously, contention for access to main memory can greatly increase a task's WCET. We introduced an analysis methodology that computes upper bounds to task delay due to memory contention. First, an arrival curve is derived for each core representing the maximum memory traffic produced by all tasks executed on it. Arrival curves are then combined with a representation of the cache behavior for the task under analysis to generate a delay bound. Based on the computed delay, we show how tasks can be feasibly scheduled according to assigned time slots on each core.

Interference-Aware Tesource Allocation for Predictable Multicore Architectures (BSC)
Inter-task interferences might appear in a multi-core when tasks try to simultaneously access the same shared hardware resources. In the past BSC has worked in processor architectures that take inter-task interference into account. In particular, the design of the architecture ensures that the maximum waiting time for an HRT task to get access to a hardware shared resource is bounded. During this period we have focused on an interference-aware allocation algorithm that considers not a single WCET estimation but a set of WCET estimations per task [PQCDV10]. Our allocation algorithm is based on two novel concepts: the WCET-matrix and the WCET-sensitivity. The former associates each WCET estimation with its corresponding execution environment. The latter measures the impact of changing the execution environment on the WCET estimation. Our allocation algorithm allows reducing the number of resources required to schedule a given taskset.

Combining Abstract Interpretation with Model Checking for Timing Analysis of Multi-core Software (Uppsala)
In this work, we study a multicore architecture where each core has a local L1 cache and all cores use a shared bus to access the off-chip memory. We use Abstract Interpretation (AI) to analyze the local cache behavior of a program running on a dedicated core. Based on the

cache analysis, we construct a Timed Automaton (TA) to model the precise timing information of the program on when to access the memory bus (i.e. when a cache miss occurs). Then we model the shared bus also using timed automata. The TA models for the bus and programs running on separated cores will be explored using the UPPAAL model checker to find the WCETs for the respective programs.
Based on the presented techniques, we have developed a tool for multicore timing analysis, which allows automatic generation of the TA models from binary code and WCET estimation for any given TA model of the shared bus. Extensive experiments have been conducted, showing that the combined approach can significantly tighten the estimations. As examples, we have studied the TDMA and FCFS buses. In both cases, the WCET bounds can be tightened by up to 240% and 82% respectively, compared with the worst-case bounds estimated based on cache misses and maximal delays for bus access.

Architectures for on-chip communication in future multi- and many-core processors (TU Braunscheig, Intel Labs) (see also Activity 6.1 Platform and MPSoC Design)
In the COMPOSE Project, the TU Braunschweig cooperated with Intel Braunschweig on new architectures for on-chip communication in future multi- and many-core processors. The goal was to develop predictable communication mechanisms and service guarantees for real-time and streaming applications. In existing embedded MPSoC, predictable communication comes at the cost of increased latencies for regular best-effort traffic, because this traffic class is treated as a "second class citizen". At the same time, a prioritization of traffic with guaranteed throughput requirements is not beneficial, because streaming applications are usually very latency-tolerant due to their predictable access patterns. TU Braunschweig has proposed and developed architectures for efficient combination of best-effort and real-time traffic. The key idea is to prioritize best-effort traffic for optimal latency, but limit its rate to retain throughput guarantees of real-time traffic. For the rate limitation, two alternatives have been developed. The first is based on distributed traffic shaping [DEK10], while the second exploits the buffer occupancy of real-time traffic to limit best-effort traffic [DE10a]. These mechanisms result in latency improvements of up to 47% compared to traditional real-time capable architectures. The throughput guarantees have been formally proven using a compositional performance analysis approach similar to the SymTA/S approach developed at TU Braunschweig.
The COMPOSE project has been successfully concluded in mid-2010. Some of its results have been presented and discussed at the Artist Desig Workshop “Mapping Applications to MPSoCs 2010” and at the Intel European Research and Innovation Conference. *(see also Section 3.4 Keynotes, Workshops, Tutorials)*

The architecture developed in the COMPOSE project also forms the basis of a many-core research platform which is currently developed by TU Braunschweig in the RECOMP project. This platform will be used to implement and evaluate mechanisms (both hardware and software) that support the simultaneous execution of applications with different safety criticalities.

http://www.ida.ing.tu-bs.de/en/research/projects/compose/
http://www.ida.ing.tu-bs.de/en/research/projects/recomp/

Predictable Fault-tolerant Communication in Distributed Embedded Systems (Linköping, DTU)
FlexRay has been widely accepted as the next generation bus protocol for automotive networks. This has led to tremendous research interest in techniques for scheduling messages on the FlexRay bus, in order to meet the hard realtime deadlines of the automotive applications. However, these techniques do not generate reliable schedules in the sense that they do not provide any performance guarantees in the presence of faults. In this work, we have proposed a framework for generating predictable, fault-tolerant message schedules on the time-triggered (static) segment of the FlexRay bus. We provide formal guarantees that the generated fault-tolerant schedules achieve the reliability goal even in the presence of transient and intermittent faults. Moreover, our technique minimizes the required number of

retransmissions of the messages in order to achieve such fault tolerant schedules, thereby, optimizing the bandwidth utilization.

Architecture and System Design: Other Issues

Time-Predictable Memory Hierarchies (TU Vienna)
We are exploring hierarchical memory architectures that simplify the WCET prediction of tasks. Instead of using cache memories for speeding up code execution, we propose to use hierarchical memories that are similar to scratchpad memories. These memories are filled by explicit prefetch operations that are executed in synchrony with program execution. The plan is to generate the instructions respectively the data that determine both the content and the timing of the operations that perform the memory transfers between the different memory levels are computed at code-generation time. So far we worked out options for the overall system and memory architecture, and design choices for explicitly controlled time-predictable hierarchical memory architecture.In a next step we plan to implement some promising options on an FPGA and evaluate them.

The PRET-C synchronous programming language for time predictability (INRIA and U. Auckland)
We have designed a new language called Precision Timed C (PRET-C), for predictable and lightweight multi-threading in C [ARG10a] [ARG10b]. PRET-C supports synchronous concurrency, preemption, and a high-level construct for logical time. In contrast to existing synchronous languages, PRET-C offers C-based shared memory communications between concurrent threads that is guaranteed to be thread safe. Due to the proposed synchronous semantics, the mapping of logical time to physical time can be achieved much more easily than with plain C, thanks to a Worst Case Reaction Time (WCRT) analyzer. Associated to the PRET-C programming language, we have developped a dedicated target architecture, called ARPRET, which combines a hardware accelerator associated to an existing softcore processor (namely Microblaze). This allows us to improve the throughput while preserving the predictability. With extensive benchmarking, we then demonstrate that ARPRET not only achieves completely predictable execution of PRET-C programs, but also improves the throughput when compared to the pure software execution of PRET-C. Our approach is also significantly more efficient in comparison to two other light-weight concurrent C variants (namely SC and Protothreads), as well as the well-known Esterel synchronous programming language.

Control/Data-Flow Predictable Reactive Processing (Kiel University)
As an evolution of the control-oriented Kiel Esterel Processor [LvH09], the Kiel Lustre Processor (KLP) [TvH09] has been designed for efficient and predictable execution of synchronous data-flow programs. Its key ideas are to use the implicit concurrency in Lustre programs to execute independent equations in parallel and to support Lustre clocks directly, in order to detect which parts of a program need to be executed in one tick. For scheduling, two different approaches have been developed, a dynamic scheduling, based on the run-time dependencies, and a priority based scheduling, with statically determined priorities. We have also developed a compiler from Lustre and Scade, which maps Lustre equations to hardware registers and computes priorities based on the data dependencies. The compiler from Scade also handles automata. http://www.informatik.uni-kiel.de/rtsys/krep/

Simulation and performance analysis (Cantabria, Politecnico di Milano, Politecnico di Torino)
The University of Cantabria, as a consequence of its participation to the Artemis Scalopes Project has improved significantly its native simulation and performance analysis technology implemented in the SCoPE tool. The main improvements are the modelling of data and

instruction caches and L2 caches, support of dynamic, frequency-voltage scaling and thermal modelling, Win32 support and improvements in the native-binary relation leading to better accuracy. As a consequence of the activity in the FP7 MultiCube project, the technology has been successfully applied to fast Design-Space Exploration.

-- The above is new material, not present in the Y2 deliverable –

3.2 Individual Publications Resulting from these Achievements

BSC

[PQCDV10] Marco Paolieri, Eduardo Quiñones, Frandisco J. Cazorla, Robert I. Davis and Mateo Valero. IA^3: An Interference Aware Allocation Algorithm for Multicore Hard Real-Time Systems. To appear in RTAS 2011 (The 17th IEEE Real-Time and Embedded Technology and Applications Symposium).

Universidad de Cantabria

[LOP10a] Patricia López Martínez, César Cuevas and José M. Drake. "Model-Driven Design of Real-time Component-Based Applications". 15th IEEE International Conference on Emerging Technologies and Factory Automation, ETFA 2010, Bilbao, September 2010, IEEE, ISBN:978-1-4244-6849-2.

[LOP10b] Patricia López Martínez, César Cuevas and José M. Drake. "RT-D&C: Deployment Specification of Real-Time Component-Based Applications". 36th Euromicro Conference on Software Engineering and Advanced Applications (SEAA 2010), Lille, France, September 2010, ISBN 978-0-7695-4170-9, pp. 147-155.

ETHZ

[SPCTC10] Andreas Schranzhofer, Rodolfo Pellizzoni, Jian-Jia Chen, Lothar Thiele, Marco Caccamo: Worst-Case Response Time Analysis of Resource Access Models in Multi-Core Systems. Proceedings of the 47th Design Automation Conference (DAC), ACM, Anaheim, California, USA, pages 332--337, June, 2010.

[SCT10] Andreas Schranzhofer, Jian-Jia Chen, Lothar Thiele: Timing Analysis for TDMA Arbitration in Resource Sharing Systems. IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS), IEEE, Stockholm, Sweden, April, 2010.

[PSCCT10] Rodolfo Pellizzoni, Andreas Schranzhofer, Jian-Jia Chen, Marco Caccamo, Lothar Thiele: Worst Case Delay Analysis for Memory Interference in Multicore Systems . Design, Automation Test in Europe Conference Exhibition (DATE), 2010, ACM, Dresden, Germany, pages 741--746, March, 2010.

Kiel University

[LvH09] Xin Li and Reinhard von Hanxleden. Multi-Threaded Reactive Programming—The Kiel Esterel Processor. IEEE Transactions on Computers, accepted 2010.

[TvH09] Claus Traulsen and Reinhard von Hanxleden. Reactive Parallel Processing for Synchronous Dataflow. In Proceedings of the 25th Symposium On Applied Computing (SAC'10), Special Track Embedded Systems: Applications, Solutions, and Techniques, Sierre, Switzerland, March 2010

Linköping

[TBEP10] B. Tanasa., U. Bordoloi., P. Eles., Z. Peng., Scheduling for Fault-Tolerant Communication on the Static Segment of FlexRay, Proceedings of 31st IEEE Real-Time Systems Symposium (RTSS10), San Diego, CA, USA, November 30-December 3, 2010.

TU Braunschweig

[DEK10] Jonas Diemer, Rolf Ernst, and Michael Kauschke, "*Efficient Throughput-Guarantees for Latency-Sensitive Networks-On-Chip*," in Proceedings of the Asia and South Pacific Design Automation Conference (ASP-DAC 2010), January 2010

[DE10a] Jonas Diemer and Rolf Ernst, "*Back Suction: Service Guarantees for Latency-Sensitive On-Chip Networks,*" in Proceedings of the 4th ACM/IEEE International Symposium on Networks-on-Chip (NOCS'10), May 2010

[SNE10] Simon Schliecker, Mircea Negrean and Rolf Ernst, "*Bounding the Shared Resource Load for the Performance Analysis of Multiprocessor Systems*," in *Proc. of Design, Automation, and Test in Europe (DATE)*, (Dresden, Germany), March 2010

TU Dortmund

[LSMM10] P. Lokuciejewski, M. Stolpe, K. Morik and P. Marwedel. *Automatic Selection of Machine Learning Models for WCET-aware Compiler Heuristic Generation*. In Proceedings of SMART '10: 4th Workshop on Statistical and Machine Learning Approaches to Architectures and Compilation, January 2010, pp. 3-17.

[PLM10] S. Plazar, P. Lokuciejewski and P. Marwedel. *WCET-driven Cache-aware Memory Content Selection*. In Proceedings of ISORC '10: 13th IEEE International Symposium on Object/Component/Service-oriented Real-time Distributed Computing, May 2010, pp. 107-114.

[LKM10] P. Lokuciejewski, T. Kelter and P. Marwedel. *Superblock-Based Source Code Optimizations for WCET Reduction*. In Proceedings of ICESS '10: 7th International Conference on Embedded Software and Systems, June 2010, pp. 1918-1925.

[FaLo10] H. Falk and P. Lokuciejewski. *A compiler framework for the reduction of worst-case execution times*. In The International Journal of Time-Critical Computing Systems (Real-Time Systems), 46(2):251-300, Springer, October 2010.

[LoMa10] P. Lokuciejewski and P. Marwedel. *Worst-Case Execution Time Aware Compilation Techniques for Real-Time Systems*. Springer, November 2010.

INRIA

[ARG10a] S. Andalam, P.S. Roop, and A. Girault. Deterministic, predictable and light-weight multithreading using PRET-C. Interactive presentation at Design Automation and Test in Europe Conference, DATE'10. Dresden, Germany, April 2010.

[ARG10b] S. Andalam, P.S. Roop, and A. Girault. Predictable multithreading of embedded applications using PRET-C. In International Conference on Formal Methods and Models for Codesign, MEMOCODE'10. Grenoble, France, July 2010.

TU Vienna

[KP2010] Raimund Kirner and Peter Puschner. *Time-Predictable Computing*. In Proc. 8th IFIP Workshop on Software Technologies for Future Embedded and Ubiquitous Systems, p. 23-34, 2010.

[KKP2010] Albrecht Kadlec, Raimund Kirner, and Peter Puschner. *Avoiding Timing Anomalies Using Code Transformations*. In Proc. 13th IEEE International Symposium on Object/Component/Service-Oriented Real-Time Distributed Computing, p. 123-132, 2010.

[CP2010] Bekim Cilku and Peter Puschner. *Towards a Time-Predictable Hierarchical Memory Architecture - Prefetching Options to be Explored*. In Proc. 13th IEEE International Symposium on Object/component/service-oriented Real-time Distributed Computing Workshops, p. 219-225, 2010.

University of Bologna

[BBB10] Carlo Bergonzini, Davide Brunelli and Luca Benini; Comparison of energy intake prediction algorithms for systems powered by photovoltaic harvesters, (2010), in: Microelectronics Journal, 41:11(766-777)

[LMRB10] Michele Lombardi, Michela Milano, Martino Ruggiero, Luca Benini: Stochastic allocation and scheduling for conditional task graphs in multi-processor systems-on-chip. J. Scheduling (SCHEDULING) 13(4):315-345 (2010).

University of Cantabria

[CPVM10] J. Castillo, H. Posadas, E. Villar, M. Martínez: "Fast Instruction Cache Modeling for Approximate Timed HW/SW Co-Simulation" 20th Great Lakes Symposium on VLSI (GLSVLSI'10), IEEE, 2010.

[CPDPSVAV11] D. Calvo, P. González, L. Diaz, H. Posadas, P. Sánchez, E. Villar, A. Acquaviva, E. Macii: "A Multi-Processing Systems-on-Chip Native Simulation Framework for Power and Thermal-Aware Design" ASP Journal on Low-Power Electronics (JOLPE): Special Issue on Low Power Design and Verification Techniques, in Press.

[PDV11] H. Posadas, L. Diaz, E. Villar:"Fast Data-Cache Modeling for Native Co-Simulation " Asia and South Pacific Design Automation Conference, ASP-DAC'11, IEEE, 2011.

[SFV11] C. Silvano, W. Fornaciari, E. Villar: "Multi-objective Design Space Exploration of Multiprocessor SoC Architectures: the MULTICUBE Approach" Springer, 2011.

Uppsala

[LNYY10] Mingsong Lv, Guan Nan, Wang Yi and Ge Yu. Combining Abstract Interpretation with Model Checking for Timing Analysis of Multicore Software. In the proc. of the 31th IEEE Real-Time Systems Symposium, November 30 - December 3, 2010, San Diego, CA, USA.

[GSYY10] Nan Guan, Martin Stigge, Wang Yi and Ge Yu. Fixed Priority Multiprocessor Scheduling: Beyond Layland and Liu's Utilization Bound. In the proc. of RTSS10 Work in Progress, November 30 - December 3, 2010, San Diego, CA, USA.

USAAR

[AMR10] S. Altmeyer, C. Maiza, and J. Reineke. *Resilience Analysis: Tightening the CRPD Bound for Set-Associative Caches*. Proceedings of the ACM SIGPLAN/SIGBED conference on Languages, compilers, and tools for embedded systems (LCTES), p. 153—162, 2010.

[GR10] D. Grund, and J. Reineke. *Precise and Efficient FIFO-Replacement Analysis Based on Static Phase Detection*. Proceedings of the 22nd Euromicro Conference on Real-Time Systems (ECRTS), p. 155—164, 2010.

[GR10] D. Grund, and J. Reineke. *Toward Precise PLRU Cache Analyis.* Proceedings of 10th International Workshop on Worst-Case Execution Time (WCET) Analysis, p. 28—39, 2010.

-- The above are new references, not present in the Y2 deliverable --

3.3 Joint Publications Resulting from these Achievements

[BCG+10] Roderick Bloem, Krishnendu Chatterjee, Karin Greimel, Thomas A. Henzinger, Barbara Jobstmann: Robustness in the Presence of Liveness. CAV 2010: 410-424

[CFGGMRT10] C. Cullmann, C. Ferdinand, G. Gebhard, D. Grund, C. Maiza, J. Reineke, B. Triquet, and R. Wilhelm. *Predictability Considerations in the Design of Multi-Core Embedded Systems*. *Proceedings of Embedded Real Time Software and* Systems (ERTSS), p. 36—42, 2010.

[GIKHV10] Arkadeb Ghosal, Daniel Iercan, Cristoph M. Kirch, Thomas A. Henzinger and Alberto L. Sangiovanni-Vincetelli, Separate compilation of hierarchical real-time programs into linear-bounded embedded machine code, Science of Computer Programming, 2010.

[ZNGSV10] H. Zeng, M. Di Natale, P. Giusto and A. Sangiovanni Vincentelli, Using Statistical Methods to Compute the Probability Distribution of Message Response Time in Controller Area Networks, *IEEE Transactions on Industrial Informatics* (TII) Volume 6,N. 4, pp.678-691,November 2010.

[IPEP10a] V. Izosimov, P. Pop, P. Eles, and Z. Peng, Synthesis of Flexible Fault-Tolerant Schedules for Embedded Systems with Soft and Hard Timing Constraints, in Design and Test Technology for Dependable Systems-on-chip, Raimund Ubar, Jaan Raik, Heinrich Theodor Vierhaus Eds., Information Science Publishing, 2010.

[IPEP10b] V. Izosimov, P. Pop, P. Eles, and Z. Peng, Value-Based Scheduling of Distributed Fault-Tolerant Real-Time Systems with Soft and Hard Timing Constraints, IEEE Workshop on Embedded Systems for Real-Time Multimedia, Scottsdale, AZ, USA, October 28-29, 2010.

[LPFM+10a] P. Lokuciejewski, S. Plazar, H. Falk, P. Marwedel and L. Thiele. *Multi-Objective Exploration of Compiler Optimizations for Real-Time Systems*. In Proceedings of ISORC '10: 13th IEEE International Symposium on Object/Component/Service-oriented Real-time Distributed Computing, May 2010, pp. 115-122.

[LPFM+10b] P. Lokuciejewski, S. Plazar, H. Falk, P. Marwedel and L. Thiele. *Approximating Pareto Optimal Compiler Optimization Sequences - a Trade-Off between WCET, ACET and Code Size*. Submitted to: Software - Practice and Experience, 2010 (Current Status: Accepted after minor revision).

-- The above are new references, not present in the Y2 deliverable --

3.4 Keynotes, Workshops, Tutorials

Key Note: **DATE 2010, Everything is Connected.**

Dresden, March 9, 2010

Alberto Sangiovanni Vincentelli gave one of the two key notes at DATE, a leading conference in design technology with more than 1,000 attendants. The talk was about the importance of

distributed systems in the world of the future, what problems we may have to face in this world and how to design complex distributed systems.

http://www.ecsi.org/date-2010-conference

Invited Talk: *Peter Marwedel, Heiko Falk:* **Reconciling compilers and timing analysis Industrial Workshop and Exhibition at the CPSWEEK 2010**
Stockholm, Sweden – April 12, 2010

Most embedded/cyber-physical systems have to respect timing constraints. Ensuring meeting such constraints is currently typically based on a trial-and-error procedure involving many time-consuming software generation attempts. In this talk, we will demonstrate how the integration of timing analysis into a compiler for an automotive processor can provide a systematic path toward optimized worst-case execution times and can cut down costs. http://www.mrtc.mdh.se/CPSweek/industrialWS/

Invited talk: *Alain Girault*: **Predictable multithreading of embedded applications using PRET-C**. Workshop on New Perspectives in Engineering and Computing for Embedded Mission Critical Systems, Thalès, Palaiseau, France, November 2010.

Key Note: Plenary Talk at the CPS week in Stockholm, Cyber Physical Systems: the Dream of Dr. Frankenstein

Stockholm, April 14, 2010

Alberto Sangiovanni Vincentelli gave one of three Plenary Talks at CPS week 2010 that hosted 5 conferences and several workshops. The talk was about forward looking applications of Cyber Physical Systems and methodology to reduce the complexity of the design.

http://www.kth.se/ees/omskolan/organisation/centra/access/dls/cpsweekplenary-1.58510?l=en_UK

Key Note: 2010 Symposium on Industrial Embedded Systems (SIES) Conference, **Connections, connections and connections. The problems of the embedded systems of the future**

Trento July 7th, 2010

Alberto Sangiovanni Vincentelli gave the key note at the Conference stressing the problems that stem from emerging behavior of widely distributed embedded systems.

http://events.unitn.it/en/sies2010

Key Note: Emerging Technologies and Factory Automation (ETFA) 2010, Distributed System Design: A Nightmare 'in fieri'

Bilbao, September 14, 2010

Alberto Sangiovanni Vincentelli gave a key note at the Conference addressing the nightmares that may ensue from the distributed system design problems.

http://www.etfa2010.org/

Key Note: IEEE System on Chip Conference (SOCC) 2010, SoC Design as an Example of Component-Based Design of Distributed Systems

Las Vegas, September 27, 2010

Alberto Sangiovanni Vincentelli gave the key note at the main conference on systems on chip outlining the need for a rigorous component-based design methodology to address the design of very large chips.

http://www.ieee-socc.org/SOCC2010/Program/program.html

Key Note: IEEE International Behavioral Modeling and Simulation Conference, Away from Plug and Pray towards Plug and Play in Analog-Mixed Signal Design: A Tale of Design Re-use

San Jose', September 24, 2010

Alberto Sangiovanni Vincentelli gave the key note at this mainly analog design conference stressing the need for compositional reasoning in analog design thus enabling a better approach to analog design re-use.

http://www.bmas-conf.org/program.html

Presentation: Predictablity in General-Purpose Many-Cores
Intel® European Research & Innovation Conference, ERIC 2010
Braunschweig, Germany, September 21-22, 2010

At the ERIC Conference, a Poster presentation was given by Jonas Diemer (TU Braunschweig) on novel network-on-chip architectures that enable future general-purpose many-core to efficiently execute real-time applications with guaranteed timing requirements. This topic was also presented at the parallel 10-year Jubilee of Intel Braunschweig, who funded the corresponding project COMPOSE.

http://www.intel.com/corporate/education/emea/event/irc/deu/

Tutorial: Architectural Aspects of Deriving Performance Guarantees
ISCA 2010
Saint Malo, France – 20.6.2010
Often, the constraints imposed by a particular application domain require a distributed implementation of embedded systems. In this case, a number of software or hardware components communicate via some interconnection network. We find this structure on various levels of granularity, starting from multiprocessor systems on a chip via distributed embedded control units (ECU) in automotive applications and ending at large-scale sensor networks.

For example, architectural concepts of heterogeneity, distributivity and parallelism can be observed on single hardware components themselves, as they are often implemented as so-called systems-on-chip (SoC) or multiprocessor-systems-on-a-chip (MPSoC). In these components, a collection of memories and heterogeneous computing resources are implemented on a single device, and communicate using networks-on-chip (NoC) that can be regarded as dedicated interconnection networks involving adapted protocols, bridges or gateways.

Embedded systems are typically reactive systems that are in continuous interaction with their physical environment to which they are connected through sensors an actuators. Examples are applications in multimedia processing, automatic control, automotive and avionics, and industrial automation. Therefore, many embedded systems must meet real-time constraints, i. e. they must react to stimuli within a time interval dictated by the environment. It becomes apparent that heterogeneous and distributed embedded real-time systems as described above are inherently difficult to design and to analyze because of the tight interaction between computation, communication and the available resources.

Part of this difficulty is caused by the fact that the functional and extra-functional behavior of the system is influenced by interferences on shared resources such as processors, memory or communication devices. Packet streams or tasks may prevent each other from using these resources, even if they belong to independent parts of the application. As a result, resource sharing strategies infuence the system behavior to a large extend. In addition, the system environment which continuously interacts with the embedded system may vary and cause an additional degree of nondeterminism.

During the system level design process of an embedded system, a designer is typically faced with questions such as whether the timing properties of a certain system design will meet the design requirements, what architectural element will act as a bottleneck, or what the on-chip memory requirements will be. Consequently it becomes one of the major challenges in the design process to analyze specific characteristics of a system design, such as end-to-end delays, buffer requirements, or throughput in an early design stage, to support making important design decisions before much time is invested in detailed implementations. This analysis is generally referred to as system level performance analysis. If the results of such an analysis is able to give guarantees on the overall system behavior, it can also be applied after the implementation phase in order to verify critical system properties.

In the presentation, we covered the following aspects of system level performance analysis of distributed embedded systems: Approaches to system-level performance analysis. Requirements in terms of accuracy, scalability, composability and modularity. Modular Performance Analysis (MPA): basic principles, methods and tool support. Examples that show the applicability: An environment to map applications onto multiprocessor platforms including specification, simulation, performance evaluation and mapping of distributed algorithms; analysis of memory access and I\O interaction on shared busses in multi-core systems.

Tutorials: Model Based System Engineering at the 2010 Control and Decision Conference (CDC)

Atlanta, Georgia – December 15th-19th, 2010

Alberto Sangiovanni Vincentelli co-organized and co-chaired with John Baras of University of Maryland two tutorial sessions at the CDC 2010 where he also presented two talks on Platform-Based Design and Model Based Design in the context of industrial applications.

Workshop: Green and Smart Embedded System Technology: Infrastructures, Methods and Tools (GREENEMBED) at the Cyber-Physical System Week

Organizing committee, general chairs, Alberto Sangiovanni Vincentelli, Huascar Espinoza, Marco Di Natale, Roberto Passerone

Stockholm, Sweden, April 12th, 2010,

Efficient production, transmission, distribution and use of energy are fundamental requirements for our modern society and the challenge of a green, low carbon economy. Embedded systems have an important role to play in increasing the energy efficiency and in reducing carbon emissions to sustainable growth. Indeed, most systems for monitoring and control of energy production, distribution and use are today interconnected and controlled by embedded devices, in areas such as industrial manufacturing, transportation systems, building automation, domestic appliances and more. This offers the opportunity for the creation of new integrated systems offering new products, processes and services with greater efficiency and better situation awareness to end-users and service and infrastructure owners.

Energy-efficient systems offer unique challenges to the embedded system community, from system-level design to dynamic and adaptive controls, optimization of architectures and communication, real-time and reliable services as well as reusable software components and systems.
Energy efficient solutions include both local and global smart solutions. Smart embedded solutions merge ubiquitous computing and the Internet of Things, i.e., the technology integration with sensors, actuators, micro-chips, micro- and nano-embedded systems that allow for collecting, filtering and producing more and more information locally, to be further consolidated and managed globally according to business functions and services. Locally, embedded systems provide information on energy consumption of every energy consuming appliance in a single location (e.g., home, building, vehicle) to be provided in real-time, in a

user friendly way, thereby empowering citizens to take decisions that lead to energy savings. Globally, energy efficient solutions include smart grid concepts, which require dynamic controls for balancing and organizing production from renewable and conventional sources, negotiating, purchasing and routing power requests, but also regulating, balancing and controlling the amount of electrical power that systems consume. From the system-level design perspective, there is a need for simulation, modelling, analysis, and monitoring methods and tools to facilitate an integrated system approach. Today, energy efficient solutions are developed by independent companies whose products or components are tested for individual performance independently of each other. An integrated system approach to the design and implementation, where these components are integrated in a way that they reduce energy consumption through cooperation, is rarely used. This often leads to significant system-level inefficiencies. System design methods and tools, including model-based solutions, must consider the growth and evolvability of hardware and software platforms, to ease the conception, development, validation and integration of new devices and services. The challenge and opportunities not only lie in the integration issue, but also in providing methods and tools for innovative solutions that satisfy government regulations, customer expectations and meet environmental challenges.

http://www.artist-embedded.org/artist/Overview,1928.html

Workshop: 8th IFIP Workshop on Software Technologies for Future Embedded and Ubiquitous Systems (SEUS 2010)
Waidhofen/Ybbs, Austria, October 13-15, 2010

The 8th IFIP Workshop on Software Technologies for Future Embedded and Ubiquitous Systems was held in Waidhofen an der Ybbs, Austria. The workshop program included a number of presentations by representatives from ArtistDesign partners. Presenters from ArtistDesign presented new predictability measures for real-time software, as well as strategies for constructing time-predictable hardware and software. A discussion on how to design temporally predictable systems for the automotive domain was one of the highlights of the workshop.
http://pan.vmars.tuwien.ac.at/seus2010/index.html

Workshop: The 8th International Conference on Formal Modeling and Analysis of Timed Systems (FORMATS 2010),
IST Austria, Klosterneuburg, Austria. 8-10 September 2010

Timing aspects of systems from a variety of computer science domains have been treated independently by different communities. Researchers interested in semantics, verification and performance analysis study models such as timed automata and timed Petri nets, the digital design community focusses on propagation and switching delays while designers of embedded controllers have to take account of the time taken by controllers to compute their responses after sampling the environment. Timing-related questions in these separate disciplines do have their particularities. However, there is a growing awareness that there are basic problems that are common to all of them. In particular, all these sub-disciplines treat systems whose behaviour depends upon combinations of logical and temporal constraints; namely, constraints on the temporal distances between occurrences of events. The aim of FORMATS is to promote the study of fundamental and practical aspects of timed systems, and to bring together researchers from different disciplines that share interests in modelling and analysis of timed systems. Typical topics include (but are not limited to):

- Foundations and Semantics: Theoretical foundations of timed systems and languages; comparison between different models (timed automata, timed Petri nets, hybrid automata, timed process algebra, max-plus algebra, probabilistic models).

- Methods and Tools: Techniques, algorithms, data structures, and software tools for analyzing timed systems and resolving temporal constraints (scheduling, worst-case execution time analysis, optimisation, model checking, testing, constraint solving, etc.).

- Applications: Adaptation and specialization of timing technology in application domains in which timing plays an important role (real-time software, hardware circuits, and problems of scheduling in manufacturing and telecommunication).

The conference was chaired by Thomas A. Henzinger and Krishnendu Chatterjee from IST Austria. http://pub.ist.ac.at/formats2010/

Summer School: UPMARC Summer School on Multicore Computing,
Uppsala, Sweden, June 21-24, 2010
The objective of the school is to offer tutorials related to parallel programming and multicore computing. This year covered lectures on predictable scheduling (by John Andersson), on predictable sharing in parallel programs (by Tobias Wrigstad), and several lectures on issues connected with distributed memory systems and their impact on programming.
http://www.it.uu.se/research/upmarc/summerschool10

Panel: 2010 Design Automation Conference, Designing the Always-Connected Car of the Future
Anaheim, California, June 15th, 2010
The panel was co-organized and chaired by Alberto Sangiovanni Vincentelli. The automotive industry is introducing novel features, such as seamless vehicle-to-vehicle and vehicle-to-infrastructure connectivity to improve in vehicle driver safety (e.g., forward collision) and comfort (e.g., routing to avoid congestion) while facing stricter government regulations, and shortened time-to-market. As a result, automotive Electronic Control System (ECS) architectures are becoming increasingly complex. To cope with these challenges and opportunities, the entire automotive supply chain is engaged as follows: automotive OEMs are managing complexity by reusing legacy components and enabling new technologies; tier one suppliers are increasingly up-integrating features on the same computing platform; tier two suppliers are providing multicore and other powerful technologies; academic institutions are doing research in new analysis, synthesis and optimization methods; and tool providers are trying to raise the level of abstraction for system modeling, analysis and optimization.
http://www.dac.com/conference+program.aspx

-- The above is new material, not present in the Y2 deliverable --

4. Overall Assessment and Vision for the Transversal Activity

4.1 Assessment for Year 3

The initiated collaborations in the context of this activity have progressed and developed further during year 3. We give some examples of resulting progress on several topics.

The PREDATOR project, which involves several partners of this activity, has entered its final year and continued the successful work of Years 1 and 2. Examples of novel collaborations include context-switch-cost-aware scheduling (USaar, AbsInt, SSSA), and clarifications of the notions of predictability (USaar, ETHZ). PREDATOR partners and IST Austria succeeded in advancing the understanding predictability on a formal basis, although this topic is far from sufficiently explored.

In Year 3, several partners, including Braunscweig, ETHZ, Linköping, and Uppsala achieved substantial progress on the problem of analyzing the predictability and intereference on shared buses and memories in multi-core systems. An interesting topic for futher research is to develop a formal measure that describes predictability and efficiency in this context. This will prove necessary to compare various architectures and resource sharing methods.

The organization of a workshop on predictability and performance, which was started at ESWEEK 2009, intending to spread issues and ideas concerning predictability, will continue with a second workshop at DATE 2011. The number of high-quality submission has increased since the previous edition. In particular, this serves well to raise the awareness on predictability issues in the industrial community. We also see that concerns of predictability are growing in industry, and the pressure to adress the problems of fast but unpredictable systems are being recognised.

Concerning the goal of reconciling timing analysis with compilation, the WCET-aware compiler WCC, developed at TU Dortmund, has matured in the past 12 months, partly in connection with strengthened collaboration between TU Dortmund, USaar, AbsInt, ETHZ, and Pisa. WCC is now able to generate and optimize industrial code, e.g., representing an engine control system, with substantially lower WCET, compared to the GCC compiler. In ArtistDesign year 3, WCC has been extended towards code generation and optimization for multi-process systems. This included porting a real-time operating system to the TriCore architecture mainly supported by WCC.

-- The above is new text, not present in the Y2 deliverable --

4.2 Overall Assesment since the start of the ArtistDesign NoE

Overall, the work has steadily progressed by considering threats to predictability, assessing and quantifying their influences, and developing solutions for redesigning architectural solutionas and/or enhancing the power analysis techniques that deliver bounds on worst-case performance. The work has grown from considering increasingly complex single-core phenomena, such as preemption and effects of software design idioms, to an increasing activity on problems related to multicore platforms. Techniques for predictable sharing of the resources on a multicore platform have already been developed, and work is intensifying on many aspects of predictable multicores.

The work on the WCET-aware compiler WCC, which previously targeted for single-core platforms has progressed also towards multi-threaded code, illustrating that this trend to consider more and more advanced features is also present in the development of mature tools.

A new trend has also emerged towards specific programming languages for time predictability, jointly with dedicated architectures. PRET-C is an example of such languages; it extends C with synchronous constructs to express multi-threading, communication with the environment, pre-emption, and logical ticks (in a manner similar to ESTEREL). Thanks to the synchronous abstraction, PRET-C provides communication through shared variables, such that communications are both deterministic and guaranteed to avoid race conditions.

-- The above is new text, not present in the Y2 deliverable --

4.3 Indicators for Integration

The Y2 report contains an account of some indicators. We reflect their development below.

- 8 joint publications (the same level as Y and Y2)
- In the last call for European project, a large number of proposals have been submitted, that build on collaboration between partners of the Predictability activity
- Educational events (summer schools) organized or co-organized by the partners.

The predictability activity organising a workshop on predictability, which will take place at DATE 2011. This workshop is meant to bring the issues of predictability in the minds of a wider audience, in particular in industry. The workshop is supported by ArtistDesign, PREDATOR and MERASA. http://ppes2011.cs.uni-saarland.de

Jointly with the workshop, it is planned to produce a joint technical paper which covers a wide selection of topics connected with timing predictability, partly based on contributions to the workshop.

As one important indicator for integration, the extensive collaboration between Dortmund, Pisa and Zurich in the context of the WCET-aware compiler WCC can be mentioned. An intensive cooperation with ArtistDesign's operating systems cluster has been set up, in particular with the cluster leader and core member SSSA-Pisa and with the affiliated member Evidence. Furthermore, Zurich's tool box of evolutionary algorithms has been integrated successfully into the WCC compiler. These integration activities led to an official port of the ERIKA operating system, which is publicly available on Evidence's websites, and to joint publications.

-- Changes wrt Y2 deliverable --

Adapted from Year 2, but extended to cover planned developments.

4.4 Long-Term Vision

In the future, predictable code generation using the WCET-aware C Compiler WCC will be extended towards multi-task and multi-core systems. For multi-task systems, it is crucial for the compiler to be aware of accesses to resources shared by the different tasks. Furthermore, detailed knowledge about the underlying operating system must be available inside the compiler. For multi-core systems, the compiler has to be aware of accesses to resources shared among the individual cores, like, e.g., shared busses or memories (including the multi-

level caches). It is envisioned that WCET-aware compilation will reduce accesses to resources shared within a single core and shared between different cores, in order to improve a system's worst-case timing. Additionally, the compiler will be able to tailor the underlying operating system, in particular its scheduling policy and its scheduling parameters. It is also envisioned the need to generate parallelization schemes that (1) improve the performance of time critical applications, and (2) eases the computation of the WCET estimation using static analysis and measurement based tools.

As a stepping stone for future research, it would be great to determine a set of indicators that can formally describe the degree of predictability of architecture components such as caches, pipelines, buses and ressource sharing strategies. They should allow for a compositional analysis and therefore, for a fair comparison of various strategies to determine predictable architectures in the future.

-- Changes wrt Y2 deliverable --

Updated wrp. to Year 2

5. Transversal Activity Participants

5.1 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 ArtistDesign	SW Synthesis and Code Generation; collaboration with TU 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	Member of the Euromicro Technical Committee on Real-Time Systems, the steering committee of the Euromicro

in the area	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	Main areas of research: Embedded Systems and Software

ARTISTDesign	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 Federal Order of Merit 2010
Further Information	Co-founder of AbsInt Angewandte Informatik GmbH Scientific Director of the Leibniz Center for Informatics Schloss Dagstuhl

5.2 Affiliated Academic Partners

	Dr. Francisco J. Cazorla (Barcelona Supercomputing Center) http://personals.ac.upc.edu/fcazorla/
Technical role(s) within ArtistDesign	Affiliated to “Design for Predictability and Performance”.
Research interests	• Multithreaded architectures and the interaction between the Operating System and the processor architecture for both high-performance and real-time systems
Role in leading conferences/journals/etc in the area	• PC member of APGES, EMSOFT, SLA++P, SRDS, TCMC. • Co-organizer of Synchronous Programming (SYNCHRON).
Notable past projects	Program committee of conferences on computer architecture as well as reviewer in many of them. I have been reviewer of: (Real-Time System Symposium) RTSS, (International Symposium on Computer Architecture) ISCA, Euromicro Conference on Digital System Design (DSD), IEEE Transaction on Embedded and computing Systems (TECS), IEEE Transactions on Computers (TC), ACM Transactions on Architecture and Code Optimization (TACO), Code Generation and Optimization (CGO), IEEE Computer Architecture Letters (CAL), High-Performance Computer Architecture (HPCA) … I have been program committee of: Architecture of Computing Systems (ARCs), HPCA, IEEE International Parallel & Distributed Processing Symposium (IPDPS), Symposium On Applied Computing (SAC), International Conference on Very Large Scale Integration (VLSI-SoC) and the Computer Architecture and Operating System co-design (CAOS) Workshop.
Awards / Decorations	• I have been selected as one of the 100 Spanish ‘leaders of the future’ according to the May 2009 issue of the Capital Magazine. This issue seeks for the 100 young Spanish citizens that will most influence Spain’s future in all innovation areas. (www.capital.es).

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, EMSOFT, 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), 2009 (2nd) and 2010 (1st)

6. Internal Reviewers for this Deliverable

- **Full Name** (affiliation)
- **Full Name** (affiliation)