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

PUBLISHABLE SUMMARY

Grant Agreement number: 214373

Project acronym: ARTISTDESIGN

Project title: ArtistDesign – Design for Embedded Systems

Funding Scheme: Full cost

Period covered: from 1st January 2008 to 31st March 2012

Name of the scientific representative of the project's co-ordinator[1], Title and Organisation:

Joseph SIFAKIS, Bruno BOUYSSOUNOUSE – UJF VERIMAG (partner n°2)

Tel: +33 4 56 52 03 51

Fax: +33 4 56 52 03 50

E-mail: Joseph.Sifakis@imag.fr, Bruno.Bouyssounouse@saxel.eu.com

Project website address: http://www.artist-embedded.org/

[1] Usually the contact person of the coordinator as specified in Art. 8.1. of the Grant Agreement.

Versions

number	comment	responsible	date
1.0	First version delivered delivered	O. Guérard	June 2012
2.0	Second, revised version	B. Bouyssounouse	Sept 2012

Table of Contents

1. Final publishable summary report

1.1 Executive summary

The ArtistDesign NoE is the visible result of the ongoing integration of a community.

ArtistDesign has been a driving force for federating the European research community in Embedded Systems Design. It brought together 31 of the best research teams as core partners, 15 Industrial and SME affiliated Industrial partners, 25 affiliated Academic partners, and 5 affiliated International Collaboration partners who participate actively in the technical meetings and events.

The central objective for ArtistDesign is to build on existing structures and links forged since 2001, to become a virtual Centre of Excellence in Embedded Systems Design. This has been achieved through tight integration between the central players of the European research community. These teams have already established a long-term vision for embedded systems in Europe, which advances the emergence of Embedded Systems as a mature discipline.

ArtistDesign built on existing international visibility and recognition, to play a leading role in structuring the area.

The research effort integrated topics, teams, and competencies, grouped into 4 Thematic Clusters: “Modelling and Validation”, “Software Synthesis, Code Generation, and Timing Analysis”, “Operating Systems and Networks”, “Platforms and MPSoC”. “Transversal Integration” covering both industrial applications and design issues aims for integration between clusters.

The NoE had a very dynamic International Collaboration programme[2], interacting at top levels with the best research centers and industrial partners in the USA: (NSF, NASA, SRI, Boeing, Honeywell, Windriver, Carnegie Mellon, Vanderbilt, Berkeley, UPenn, UNC Chapel Hill, UIUC, etc) and in Asia (Tsinghua University, Chinese Academy of Sciences, Seoul National University, East China Normal University, etc).

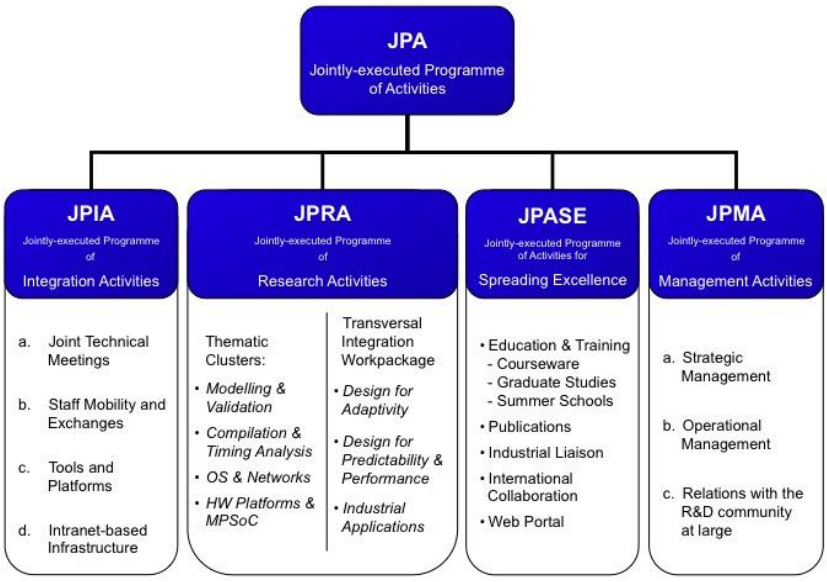

ArtistDesign also had a very strong tradition of Summer Schools and Graduate Schools[3], and major workshops[4].

ArtistDesign has built on existing international visibility and recognition, and played a leading role in structuring the area. It has provided a significant evolutionary step for integrating the leading embedded systems design research teams - both in terms of impact on the overall structuring and lasting integration within the consortium and more generally within the area in Europe.

[2] http://www.artist-embedded.org/artist/-International-Collaboration,1050-.html
[3] http://www.artist-embedded.org/artist/-Schools-.html
[4] http://www.artist-embedded.org/artist/-Workshops-and-Seminars,29-.html

1.2 Summary description of project context and objectives

1.2.1 Structure and Vision for the Research Effort

The NoE was composed of intra and inter-cluster research activities on cutting-edge topics in embedded systems design. While the main bulk of financing for these activities was taken up by outside programmes (Integrated Projects, National Programmes, Industrial Contracts, etc), the ArtistDesign NoE financed the extra effort derived from integrating these into a single coherent research programme.

Thus, the essential ingredient within ArtistDesign was the JPRA, which motivated the participating research teams far more than the actual financing, was tiny in comparison with the overall research aims. It was completed by the Joint Programme of Integrating Activities (JPIA), and the Joint Programme of Activities for Spreading Excellence (JPASE), and overseen by the Joint Programme of Management Activities (JPMA).

The structure for the research activities reflects the following decomposition of the embedded systems design flow.

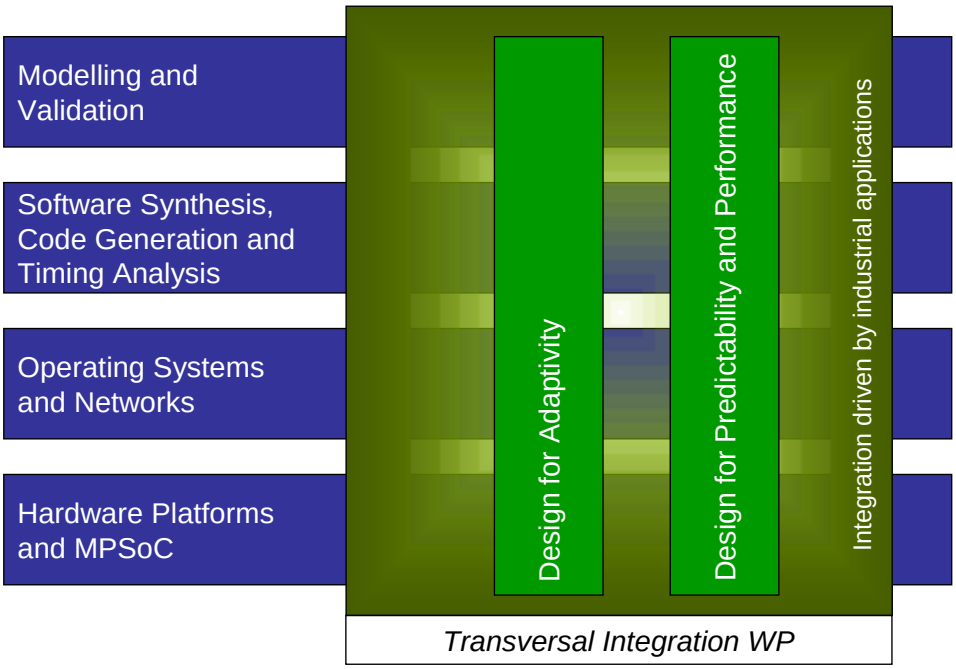

This design flow was composed of the following cooperating activities, starting with component-based modelling and leading to implementation. These activities must be well coordinated, and supported by tools and methods to ensure satisfactory levels of productivity and quality. Accordingly, we have structured the area of embedded systems design into the following topics.

1.2.2 Joint Programme of Research Activities (JPRA): Thematic Clusters

- Modelling and Validation.
  Unlike other computer systems, embedded systems are strongly connected with a physical environment. A scientific foundation for embedded systems must therefore deal simultaneously with software, hardware resources, and the physical environment, in a quantitative manner. In order to gain independence from a particular target

platform, embedded system design must be model-based. In order to scale to complex applications, embedded system design must be component-based. The overall objective of this activity was to develop model and component-based theories, methods, and tools that establish a coherent family of design flows spanning the areas of computer science, control, and hardware. The activity brings together the most important teams in the area of model and component based design in Europe.

- SW Synthesis, Code Generation and Timing Analysis.
  There was a continuing demand for higher performance of information processing, which stimulates using a growing amount of parallelism (including using multiple processors). This trend affects the design of embedded systems. We address issues related to multiple heterogeneous processors on a chip, also containing memory hierarchies and communication interfaces. Such processors can only be exploited if (sets of) applications can be efficiently mapped to heterogeneous processors. Timing analysis was also affected by the trend toward the new platforms. Timing analysis has to cope with the kind of memory hierarchies found in MPSoCs. Also, timing analysis beyond the single processor was required. Hence, timing analysis will also consider the timing of communication. The overall objective was to provide safe timing guarantees for systems consisting of local memories hierarchies and multiple processors.

- Operating Systems and Networks.
  We investigate how current real-time operating systems have to be extended or modified to support emerging real-time embedded systems characterized by a high degree of complexity, highly variable resource requirements and parallel processing such as multicores. Most embedded systems are often characterized by scarce resources in terms of processing power, memory, space, weight, energy, and cost. Hence, another objective was to investigate kernel mechanisms that can efficiently manage the available resources, taking multiple constraints into account, whilst guaranteeing isolation properties. Also, to support dynamic applications with variable resource requirements or to cope with unpredictable resource availability, feedback control techniques for resource management at the operating system and application level are also investigated.

- Hardware Platforms and MPSoC Design.
  While hardware platforms for embedded applications will continue to be multi-core, with increasing degrees of parallelism, the evolution trajectory on programming models, design-time and run-time application environments was much less clear. The consequence was fragmentation: while many research teams are working on one or more of these domains, there was little communication and integration, this leads to duplication of results and overall slow progress. The teams involved in this activity have a wide-ranging research experience, which covers all the key areas in MPSoC application specification mapping. The integration activity supported by ArtistDesign will help the participants to the cluster in strengthening the coherency of their approaches and focus on addressing complementary issues in a synergistic fashion.

1.2.3 Joint Programme of Research Activities (JPRA): Transversal Integration Activities

- Design for Adaptivity.
  An embedded hardware-software system was adaptive, if it can modify its behaviour and/or architecture to changing requirements. Adaptivity was increasingly important as the complexity and autonomy of embedded systems increases. Adaptivity was a cross-cutting system characteristic that affects both hardware and software. At the software-level adaptivity was mainly concerned with flexible and adaptive resource scheduling, e.g., CPU time scheduling. At the hardware-level adaptivity includes both adaptation of operation modes, e.g., supply voltage and clock frequency, processor instruction sets,

and dynamic management of hardware resources, e.g., processing elements and memory.

- Design for Predictability and Performance.
  Many applications have strict requirements on timing, and limited resources (memory, processing power, power consumption, etc.). All systems also have increasing demands on (average) performance, which has motivated the introduction of features such as caching, pipelining, and (now becoming very prominent) multiprocessor platforms. Almost all such efficiency-increasing features drastically increase variability and decrease analyzability of response-times, etc. and thus have a detrimental effect on predictability. Since the introduction of new architectural features was inevitable, it was important to: a) develop technology and design techniques for achieving predictability of systems built on modern platforms, and b) investigate the trade-offs between performance and predictability.

- Integration Driven by Industrial Applications.
  To have a strong impact on industry and society at large, the results of the Thematic Clusters need to be harmonized in an overall design flow that can sustain the embedded design chain from conception of the product to its implementation. The design chains vary in length and players according to the industrial segment addressed: for example, the design chain in automotive electronics starts with the car maker (e.g., BMW, Daimler Chrysler, Peugeot, Fiat), goes through the Tier 1 suppliers (e.g., Contiteves, Bosch, Magneti Marelli) and connects to the Tier 2 suppliers (e.g., FreeScale, ST, Infineon, Hitachi). It often includes IP providers such as programmable cores, RTOS and software development tool providers and design service companies. In the mobile communication domain, the chain starts with the application developers (e.g., gaming and video content), includes the telecommunication operators (e.g., Telecom Italia and Telefonica), the device makers (e.g., Nokia and Ericsson), the silicon makers (e.g., TI, Qualcomm and ST) and outsourcing manufacturing companies (e.g., Flextronics). Today, there was stress in the chain as the technology advances may create opportunities to redefine the roles of the various players. In addition, system integrators are often faced with an almost impossible task of composing their design out of parts supplied by companies whose design methods and standards are widely different and about which they have limited or no information. There was a need for an all-encompassing approach to system design that can make an entire industrial segment work as a virtual vertically integrated company.

1.2.4 Joint Programme of Activities for Spreading Excellence (JPASE)

The JPIA activities promote integration of geographically dispersed teams and have long-lasting effects:

- Joint Technical Meetings.
  Joint Technical meetings aim to present, discuss and integrate the ongoing work, and exchange points of view with other teams. They also serve to identify future work directions.

- Staff Mobility and Exchanges.
  This is essential for integration within the NoE, including mobility of students and/or researchers, between core teams, or between core teams and affiliated teams. Mobility is justified by and refers to involvement in an activity from the JPRA or JPIA, or one of the following: co-funded scholarships with industry; exchange of students and personnel within the consortium.

- Tools and Platforms.
  A research platform is composed of competencies, resources, and tools targeting specific technical and scientific objectives around a chosen topic. These are at the state-of-the-art, and are made available to the R&D community for experimentation, demonstration, evaluation, and teaching.

  The research platforms, tools and facilities are an essential tool for implementing the JPIA. They will lay the groundwork for the JPRA, allowing common research to occur and capitalisation on research results. Platforms are used as the basis for transfer of research results to industry. They allow teaching practical knowledge of the concepts and techniques.

  ArtistDesign platforms are not defined from scratch – they integrate the results of long-term efforts, and are meant to be durable, evolving with the state of the art. The partners are committed to durability, and have invested significant resources into their development. The construction of ArtistDesign has provided the opportunity to assemble existing pieces into a rationally-structured set of platforms, covering the area of embedded systems design.

  Some of the ArtistDesign platforms have international visibility, and the ambition is for these to serve as world-wide references in their respective topics.

1.2.5 Joint Programme of Activities for Spreading Excellence (JPASE)

The NoE's actions for Spreading Excellence have been at 3 levels:

- Targeted towards affiliated partners
  Affiliated partners are not core members in the consortium, but receive support for travelling to ArtistDesign meetings, and actively contribute to the implementation of the Joint Programme of Activities (JPA). These affiliated partners include industrial, SME, academic, and international collaboration affiliates.
- Targeted towards the scientific and technical community in the large
  This is achieved mainly bottom-up through the organisation of scientific events, publications, distribution of tools and components, industrial partnerships (not funded by ArtistDesign), education; and through the ArtistDesign web pages.
- Targeted towards students
  A particular focus has been placed on the ARTIST Summer Schools this year – with a truly outstanding programme of lecturers, and the innovation of providing the lectures in video form on the Artist website.

Regarding Scientific events, we distinguish between conferences and workshops, schools, and high-level events mainly for International Collaboration.

High Level Events for International Collaboration
High-level Events are intended to gather together the very best world-leading experts from academia and industry, to discuss progress on the state of the art, relevant work directions.

Three ArtistDesign members are on the steering board for the ARTEMIS European Technology Platform. In this capacity, they participate in working groups for defining the overall European long term strategy in the area.

Publications
The ArtistDesign community has pusrsued a very active publishing policy, with a strong presence in scientific journals and conferences, as attested by the extensive list of publications provided in this document. Publication of research is a bottom-up process, which may seem

chaotic – but this is intrinsic to research.

Tools and Components

The ArtistDesign community plays a leading role in the distribution of software tools and components, on verification/validation tools. Some tools are distributed free of charge, such as UPAAL, IF. Others are commercialised, such as AbsInt, SymTA/S. For many other tools used in the platforms, and shared between the Artist partners, a common dissemination policy has not yet been defined.

Industrial Liaison

ArtistDesign has a wide array of affiliated industrial and SME partners (see the Thematic and Transversal Activity deliverables). Most of these partners have participated in some way in the ArtistDesign technical meetings and the overall effort. There is strong, high-level industry participation through the various Spreading Excellence events organised by ArtistDesign. Our active involvement in the European Technology Platform ARTEMIS also could have a significant and long-term impact.

We believe that the strong involvement of four main ArtistDesign partners in the SPEEDS Integrated Project has a very positive impact on progress in the state of the art, in component-based embedded systems engineering.

ArtistDesign leveraged on the worldwide visibility of its activities. It is progressively creating a European embedded systems design community and spreading the “Artist culture” in all major research institutions.

To ensure that the next generation of researchers will continue in this direction we, as a consortium, have devoted a great deal of effort to Spreading Excellence, in both academic and industrial circles. Furthermore, through our links with both core and affiliated partners, we actively set up permanent links between industry and public research, based on existing partner collaborations with major industrial players in the area.

The JPASE activities spread excellence and structure the community at large. They are planned by the Strategic Management Board, and are implemented by ArtistDesign core and affiliated partners.

The NoE has leveraged on its members and teams, who play a main role in the organisation of world-class scientific events, to disseminate results in the area. We expect that the NoE’s structured and authoritative dissemination will have a strong effect on the community as a whole, for orienting and creating synergy for research.

1.3 Description of the main S&T results / foregrounds

The main Scientific and Technical results are organized according to the cluster /activity structure described earlier.

1.3.1 Thematic Cluster: Modelling and Validation

Year 1

- Modeling

  On modelling heterogeneous systems, we have obtained some very significant results:

  - Work by INRIA, PARADES, and VERIMAG, on the theory of tagged systems provides conditions for correct implementation of synchronous systems on “less synchronous” architectures. Work by INRIA on new models of computation, the Kahn-extended Event Graph (KEG), which adds “static control” in the Model of Computation of Marked Graphs.
  - Work on the BIP component framework, introduces a notion of expressiveness for component-based formalisms, which provides a basis for their comparison. This notion drastically differs from the usual one, as it takes into account the expressive power of composition operators (glue operators).
  - Significant progress has been achieved in methods for distributed implementation of non-distributed specifications. VERIMAG studied a method for the automatic generation of distributed implementations of BIP models. INRIA, PARADES and VERIMAG studied the concept of loosely time-triggered architectures, implementing time-triggered architectures. Finally, fully asynchronous implementations of synchronous systems have been studied.
  - Other results on distributed implementations include reliability, new heuristics in scheduling for reliability, design of communication architectures and time-triggered system-on-chip architectures.

- Validation

  At the crossroads between Modeling and Validation, we have obtained significant results on Interfaces and Composability, including:

  - The development of interface theories supporting component reuse (EPFL). We have shown that existing interface theories provide no formal support for component reuse. We enriched interface theories with a new operation allowing the same component to implement several different interfaces in a design.
  - The development of contract-based verification techniques for the heterogeneous rich component (HRC) model (INRIA, PARADES, VERIMAG), in the framework of the SPEEDS project. The techniques allow handling multiparty interaction, as well as many different languages for describing notions of refinement under contexts. These results found application in the verification and analysis of HRC models.
  - In joint work, ETHZ and Uppsala propose modular performance analysis techniques, based on the real-time calculus and timed automata. A prototype tool name CATS for compositional timing and performance analysis has been developed.

- VERIMAG has continued the work on compositional deadlock verification of BIP programs, and its implementation in the DeadlockFinder tool. Other results include the enhancement of existing component models, such as the synthesis of controllers from specifications and the generation of component models from their observed behaviour.

On qualitative validation, work has been carried out in the following directions:

- Significant contributions to game-theoretic approaches to real-time system testing. By modelling the systems as timed game automata and specifying the test purposes as formulas, we developed a timed game solver Tiga to synthesize testing strategies.
- We studied games for different extensions of timed automata such as weighted timed automata, priced timed automata, multi-priced timed automata. The results relate to the complexity of decision problems for these automata, as well as model checking and synthesis algorithms. The notion of timed parity games has been studied, with a focus on robustness and complexity. We have also studied reachability in timed games.
- Continuing on work from previous years, we have extended and improved the functionality of the UPPAAL tool, including the use of slicing techniques for model optimization as well as features supporting interface theory for real-time systems.
- We have studied quantitative testing techniques. In particular, we have developed a theory allowing testing of systems in the presence of measurement imprecisions. We also studied testing methods for probabilistic processes.
- We have studied quantitative model checking techniques for timed models, including timed automata, linear hybrid automata and general non-linear hybrid systems. The work on verification has been applied to non-trivial case studies and systems, in particular in collaboration with industry.

Finally, we have studied compositional synthesis and verification techniques. These include modular supervisory control, as well as the verification of component-based systems.

Year 2

The principal Y2 results obtained are:

- Modeling
  - The development of modelling languages such as MARTE and Y-charts, as well as their application to non-trivial case studies.
  - The foundations of a contract-based theory for components, as well as a theory for relational interfaces. These can be used for compositional verification of embedded systems.
  - The development of model-based design methodologies for
    a) flight control systems,
    b) complex systems ranging from automobiles, buildings, and airplanes,
    c) multi-core systems.
  - The development of resource modelling techniques based on languages and models that explicitly represent different resource types, such as processors, memories and energy.

- The study of a sound semantic basis for quantitative modelling, including
  a) performance models,
  b) quantitative generalisation of classical languages,
  c) synthesis of optimal controllers from quantitative high-level specifications.

- Validation

  - The development of compositional validation techniques, including:
    a) compositional verification of deadlock freedom,
    b) modular analysis for timed systems,
    c) compositional safety analysis,
    d) compositional verification of probabilistic systems.
  - The development of quantitative validation techniques, including:
    a) the analysis of energy-related properties of sensor-networks,
    b) the verification of hybrid systems
    c) the study of multi-processor scheduling techniques and cache policies for ensuring timing predictability,
    d) the study of quantitative models and their associated verification techniques, such as timed automata, and stochastic automata,
    e) the verification of transactional memories, by model checking.
  - The development of techniques for cross-layer validation, such as:
    a) control under partial observation,
    b) adapting abstraction techniques to black-box analysis and learning,
    c) test-case generation techniques, under partial observability.

Year 3

Both research activities with the cluster – the *Modeling Activity* and the *Validation Activity* – have progressed substantially within the third year, and with significant synergy between proposed modeling formalisms and methods and validation techniques they support:

Within the sub-activity *Component Modeling* and *Compositional Validation* several partners have worked substantially and collaboratively on compositional design and verification methodologies for functional, timing and stochastic aspect. The results include

- Assume/guarantee reasoning, interface automata as well as modal transition systems for rich models.
- Theoretical foundations and coordination languages have been developed for heterogeneous systems.
- A framework for tool integration based on meta-models and model-transformations.

The work in the sub-activity *Resource Modeling (*of the *Modeling Activity)* includes:

- design space exploration,
- multi-core scheduling,
- modelling paradigms for quantitative resources
- platform models including transactional memory.

The work in the sub-activity *Quantitative Modeling (*of the *Modeling Activity)* has produced significant results on:

- design frameworks for quantitative modeling, in particular weighted automata, priced timed automata and quantitative communication models.
- synthesis of models guaranteeing quantitative properties.

Within the sub-activity *Quantitative Validation* substantial work has been made on improved schedulability analyses supporting multiprocessor and multi-core applications. The work:

- takes into account scheduler overhead for power-awareness – i.e. exploiting slacks in the system of processes to reduce power consumption while insuring deadlines are met.
- combines abstract interpretation and model-checking for timing and interference analysis of parallel programs on multi-core,
- has been applied for schedulability analysis of Safety Critical Java applications.

Within the sub-activity *Cross-Layer Validation* substantial work has been made on improved methods for model-based testing including:

- incremental testing of composite systems,
- off-line test generation from timed automata models,
- model-based test generation for data-intensive systems, as well as runtime monitoring.

Year 4

Both research activities with the cluster – the Modelling Activity and the Validation Activity – have progressed substantially within the fourth year, and with significant synergy between proposed modelling formalisms and methods and validation techniques they support:

The work on Component Modeling and Compositional Validation involved several partners that produced significant results on compositional modelling and verification:

Results on modelling can be summarized as follows:

- Composition frameworks for behaviour and properties of heterogeneous systems such as assume/guarantee reasoning, interface automata, modal transition systems as well composition frameworks for tool integration based on meta-models and model-transformations have been consolidated and applied to case studies.
- Resource modelling techniques applied to design space exploration, multi-core scheduling, performance evaluation and derivation of distributed implementations from global specifications.
- Quantitative modelling techniques for weighted automata, priced timed automata and quantitative communication models.

Results on validation can be summarized as follows:

- Quantitative Validation covering a wide range of techniques for WCET analysis, schedulability analysis, frequency analysis of timed automata, analysis of parametric quantitative models, and analysis of resource consumption using energy- and price-extensions of timed automata. These techniques use new notions of metrics and robustness.
- Cross-Layer Validation focusing on model-based testing techniques such conformance testing of real-time systems using time- and data abstractions, asynchronous testing

and test-case generation for embedded Simulink, incremental testing of composite systems as well as runtime monitoring.

In addition to these results, the Cluster has endeavoured a considerable integration effort for connecting tools, joint meetings, open workshops and joint publications.

1.3.2 Thematic Cluster: Software Synthesis, Code Generation and Timing Analysis

Year 1

- Software Synthesis and Code Generation
  - We studied the influence of scratchpad memory allocation techniques on worst case execution times (WCET). We developed integer linear programming models to decide which parts of a program's code or data can be moved onto the highly predictable scratchpad. First experiments show WCET reductions of more than 50% for several benchmarks.
    We also investigated WCET-aware register allocation techniques, by extending existing techniques based on graph colouring.
  - We continued work on scalable source-level analysis and annotation-based timing analysis methods. The SATIrE infrastructure allows building analysers that take source code annotations as additional input, and generate output as annotations. This allows a significant increase in productivity, by requiring the user to annotate the relevant timing information that cannot be automatically computed. The integration of PAG was instrumental in investigating the scalability of analyses.
  - We studied polyhedral loop parallelisation techniques for multi-core systems.
- Regarding tools and platforms, we developed work in the following directions:
  - We designed a Static Loop Analyzer, allowing to estimate loop iteration bounds. This information is essential for a large number program analyses. Our analyser improves analysis techniques based on conventional abstract interpretation by integrating a new static polytope-based loop evaluation method.
    We have demonstrated the applicability of the analyser on benchmarks taken from the benchmark suites MRTC, DSPStone, MiBench, UTDSP and MediaBench. Our loop analyser was the only tool able to answer all questions related to flow fact during the WCET tool challenge 2008.
  - In continuation of work performed in Year 3, we developed a new WCET-aware procedure positioning and cloning technique. The compiler optimisations obtained were exploited for WCET reduction. Results on real-world benchmarks show WCET reductions of 10% on average, while ACET is reduced by 2 on average.
  - The cooperation between ACE and Aachen on the retargetable code optimizations has been continued. The conditional execution engines have been extended by a strong retargeting formalism.
- Timing Analysis
  - In addition to experimental work, we developed important results on Timing Analysis.

- We have studied a notion of time predictability of cache architectures, which is the first precise notion found in the literature. Four different cache replacement strategies were compared and the LRU strategy was found to be optimal. This research is related to work within the PREDATOR FP7 project, which attempts to reconcile performance and predictability.
- The study of Timing anomalies, where local worst-case choices may not lead to the global worst-case scenario, is essential for time predictability. We have studied techniques for handling timing anomalies for efficient WCET analysis, as well as for measuring the impact of timing anomalies on WCET analysis.
- Other work on Timing Analysis includes parametric Timing Analysis, where some parameters of the program can remain unknown until execution. We also developed Timing Analysis techniques, in collaboration with BOSCH, taking into account operating modes of programs, computed semi-automatically. Finally, we developed WCET analysis for systems with preemptive scheduling.
- Work on the AIR format has continued. The format was extended and adapted to the needs of the partners. The attribute database was extended with new attributes.
- We also worked on the development and improvement of formats for ensuring the interoperability of the tools. The work on formats includes ALF for computation semantics representation, conversion of the ABSINT AIR format to SWEET format, and the definition of common flow description attributes.
- As was the case last year, the WCET Challenge 2008 consisted of a set of benchmark programs and analysis tasks to be performed by the contestants. http://www.artist-embedded.org/artist/-WCET-08-.html

Year 2

The principal Y2 results obtained are:

- Timing Analysis
  - Using model identification and learning techniques for estimating WCET.
  - Theory for ensuring timing predictability, as well as for timing composability.
  - The development of a common annotation language for WCET analysis.
  - The integration of Timing Analysis and Compilation, to achieve WCET-aware compilation.

- Software Synthesis and Code Generation
  - Techniques for mapping applications to MPSoCs, so as to meet a given set of non-functional properties.
  - Code optimization techniques, taking into account WCET,
  - Techniques for verifying code generation methods.

Year 3

In year 3, we have seen a proliferation of the basic techniques studied by this cluster. The importance of using multi-processor systems has been growing even more than it did in year 2.

Several tools for mapping of applications to MPSoCs have become available (e.g. from RWTH Aachen, IMEC, Erlangen-Nuremberg and Seoul National University). We have reached a situation where such tools can be considered state of the art. We expect such tools to leave the research labs in the not too distant future.

Within the cluster, timing analysis and timing predictability for multicore platforms have seen substantial progress. The results include

- the worst-case execution time aware compiler WCC (Dortmund, AbsInt, USAAR)
- cache-aware scheduling (USAAR, SSSA). Contacts with the MPSoC design cluster have been strengthened.
- timing analysis of multicore systems with shared caches, and to bound the context switch penalty due to cache effects in preemptive systems
- analysis on micro-architectural level has progressed, especially regarding cache replacement policies and pipeline behaviour.
- automated derivation of timing models from VHDL specifications.
- generation of timing models from observations, based on machine learning and model identification.

Concerning the goal of reconciling timing analysis with compilation, the WCET-aware compiler WCC developed at TU Dortmund has been extended beyond the initial TriCore hardware platform and toward multi-objective optimization. Collaboration between TU Dortmund, AbsInt and the ArtistDesign Cluster on Operating Systems and Networks has been strengthened. We established a new link between reliability, compilers, operating systems and real-time systems.

Finally, dissemination comprises the inclusion of educational material on software synthesis, compilers and timing analysis in the second edition of the textbook on embedded systems by P. Marwedel.

Year 4

In year 4, we have seen a further proliferation of the basic techniques studied by this cluster. The importance of using multi-processor systems has been continuing to grow. Any session on programming multi-cores and multi-processor systems is filled with people. Fortunately, ArtistDesign is active in this area.

The work on software synthesis and code generation focused on the development of tools and resource-aware compilation. We developed two tools for mapping applications to multi-core or multi-processor platforms (RWTH Aachen, IMEC). Our work on resource-aware compilation has continued with new results on energy efficiency and thermal behavior control as well as with fundamental machine-learning techniques for optimized code generation.

In program flow analysis, MDH and Tidorum have made advances towards increased soundness by developing an advanced relational value analysis that takes possible overflows and wraparounds into account. This is important for small embedded systems, where wraparounds are common

Additional activities include the organization of an international workshop on Software Synthesis (http://www.artist-embedded.org/artist/-WSS-11-.html) and development of new educational material on software synthesis, compilers and timing analysis in the second edition of the textbook on embedded systems by P. Marwedel.

The work on timing analysis and timing predictability has progressed significantly in two directions. The first focuses on enforcing predictability through determinism. It produced new and industrially relevant results on cache analysis and ache-aware memory allocation that have been taken up by commercial tools such as aiT from Absint. The second takes a probabilistic approach and relies on randomization to make timings on micro-level independent. Very promising initial results has been obtained.

Advances in hybrid WCET analysis methods, which include elements of measurements and testing, have been made (MDH, York, TU Vienna). Such timing models can be used to provide worst-case timing estimates early as well as small but appropriate sets of test vectors for tasks with very large input sets, and evaluation of coverage metrics for test-data generation.

Finally, the Cluster has achieved increased integration of timing analysis tools and compilation tools (TU Dortmund, TU Vienna).

1.3.3 Thematic Cluster: Operating Systems and Networks

Year 1

- Resource-aware Operating Systems

  In addition to work done in Artist2, we developed work in the following directions:

  - Modeling and analysis of control-driven tasks. The standard design of control is based on the periodic sampling. We studied a model which saves a considerable amount of computational resources which samples the inputs when needed.
  - Implementation of a flexible scheduling framework called FRSH that is capable of handling multiple concurrent activites with different criticality and timing in the same system. The framework has been designed to be implemented on different platforms.
  - ERIKA support for the EasyBee radio transceiver has been developed.
  - We studied issues relating to the operating system support needed by advanced users of a real-time specification for Java. In particular, two issues have been addresses: how to handle systems that contain a large number of events; and how to measure blocking time.

- Scheduling and Resource Management

  We have had a large volume of activity on this topic:

  - All the partners, under the leadership of York, have worked for establishing a taxonomy of resource usage. The taxonomy distinguishes between different classes of resources - each class being subdivided into a number of resource types.
  - The architectural model of a Flexible Scheduling Framework developed in the FRESCOR and FIRST EU-IST projects has been extended to include a contract model. Contracts represent complex requirements of the applications which can be managed by the underlying system to provide the required level of service.
  - Several activities on scheduling, in particular multi-resource scheduling for multi-core platforms, schedulability for CAN-based control applications, sensitivity analysis, flexible scheduling on low-cost microcontrollers.

- Other work related to the Transversal Activity: “Design for Adaptivity” has been carried out, including dynamic runtime adaptability, optimal period selection and scheduling for embedded controllers.

- Real-Time Networks

  We have carried out work on:

  - Analysis techniques, including Worst Case analysis and dimensioning of cluster-tree Wireless Sensor Networks, as well as analysis for specific networks.
  - We studied techniques for supporting real-time communication and QoS for Wireless Sensor Networks. These include work around the use of IEEE 802.15.4 and ZigBee as federating communication protocols for Wireless Sensor Network applications, as well as supporting real-time communication of the Erika real-time operating system.
  - We have also furthered work, started in Artist2, on student design competitions in the scope of the IEEE Real Time Systems symposium.

Year 2

The principal Y2 results obtained are:

- Scheduling and Resource Management
  - resource management, including:
    a) the development of a taxonomy of resource management as a wiki,
    b) memory resource management techniques,
    c) adaptive resource management.
  - scheduling, including:
    a) hierarchical scheduling,
    b) scheduling and placement algorithms for multi-processor systems,
    c) multi-resource contract-based scheduling,
    d) data-flow scheduling, using constraint programming.
  - language support for programming real-time systems,
  - sporadic event-based control.

- Real-Time Networks
  - Timeliness in wireless sensor networks,
  - Mobility issues in ad-hoc real-time wireless communication
  - Robust communication with star topologies,
  - Real-time support to middleware and composability,
  - Applications of wireless networks in industrial environments, in intelligent transportation systems, and in health applications.

Year 3

The fruitful collaboration among the cluster participants is demonstrated by the number of joint publications, projects and events organized within the cluster. The main examples are the

organized workshops and conferences, graduate courses, and the various research consortia that have led to new European projects, like ACTORS, PREDATOR, IRMOS, and SOOS.

All research activities in the cluster have progressed substantially within the third year. The following list briefly summarizes some of the major achievements for year 3. Details and more information can be found in the three activity reports by the cluster.

- UNIBO-PISA continued to collaborate on predictability and modularity of MPSoC for Real-Time applications. The interaction has been realized by the integrating tools developed by the partners: the Erika RTOS from SSSA has been extensively used on the MPARM platform developed by UNIBO.
- USAAR, PISA, Dortmund, AbsInt. Also supported be the PREDATOR project, these partners collaborated to improve the estimation of worst-case execution times considering cache-aware scheduling and WCET-aware compilers.
- EVIDENCE-PISA. A great effort has been done to introduce resource reservation and deadline-based scheduling (EDF) in the Linux operating system, so enabling the implementation of advanced resource reservations techniques.
- LUND-TUKL-PISA. Also supported by the ACTORS project, these partners collaborated to develop a design framework for partitioning real-time applications on multicore heterogeneous systems, with the objective of guaranteeing optimal usage of the available resources.
- PISA-UPC tightly worked together to define a laboratory platform and experiment to be integrated in the education of embedded control system engineers. The experiment consists in the control of a dynamical system on a platform supported by the Erika real-time kernel. The set up has been tested on a graduate course jointly organized in Pisa on June 14-18, 2010.
- ULUND-PISA continued to collaborate on event-based control systems. In this third year, the work has focused on network scheduling of event-based controllers.
- Aveiro, UnivPorto and Malardalen worked on a reconfigurable hierarchical scheduling framework within an enhanced Ethernet switch that allows an efficient use of bandwidth, enforcing temporal and spatial isolation.
- York, Cantabria, Porto, Madrid, Valencia collaborated for providing a language support for programming schedulable systems. This year the work has focused on getting support for multiprocessors into the next versions of Ada and the Real-Time Specification for Java. These have now effectively been agreed and will enter into the standards at their next releases.
- TUKL, CSEM, Philips, Pisa, York, Porto, Prague. Contributed on the development of timeliness in Wireless Sensor Networks. The teams at TUKL, CSEM and Philips proposed a generalized notion of timeliness, which suits the characteristics of WSN, based on the requirements in the EU IST project WASP.
- Mallorca, UnivPorto, Catalonia, IFP addresses the problem of robustness and timeliness in Controller Area Networks.
- Cantabria, Madrid: UPM and UC3M, Bilbao, UnivPorto collaborated for providing real-time support to middleware and composability. A set of timing analysis tools has been integrated with a toolset for MDE. In addition, a new approach has been explored to integrate the real-time end-to-end flow model with the automatic generation of Ravenscar-compliant source code in distribution middleware.

- ALL PARTNERS contributed to a major activity (coordinated by YORK) for building a taxonomy of resources, considering multi-resource platforms and including the use of banded notions of time and hierarchical structures.
- Madrid, Pisa, Aveiro, UnivPorto, Malardalen, NXP, TUKL worked on protocol optimizations for embedded real-time communications. The validation showed performance improvements in comparison to currently used infrastructures. The performance has been reported to a journal in an article which now in accepted status.
- Catania, Pisa, Evidence have been involved in intelligent transportation systems, automatic traffic monitoring and road surveillance. Various sensors have been used to estimate traffic parameters. Catania proposed a wireless sensor network architecture based on computer vision techniques for automatic scene analysis and interpretation.

Year 4

The work developed by the cluster involved several partners that produced significant results summarized as follows:

The work on operating systems and middleware focuses on resource reservation and predictability. We developed an implementation of a real-time scheduler in the Linux kernel, with a support for resource reservation. We also developed a programming framework to support resource reservation of concurrent real-time applications on multi-core platforms, considered by Ericsson for software development in next generation cell phones. Finally, we proposed a comprehensive taxonomy for the resources currently used in embedded real-time systems.

Our work on predictability includes cache-aware analysis and scheduling for safety-critical applications, In collaboration with the Cluster on Compilers and Timing analysis.

The Cluster also developed a middleware and communication protocol for teams of mobile robots that are self-reconfigurable and provide efficient support to intensive interactions and which have been adopted by several teams in the RoboCup Middle Size League.

The work on networks includes two toolsets. One for the design, analysis, configuration and deployment of dense WSNs. The other is the MAST (Modelling and Analysis Suite for Real-Time Applications), which was enhanced with more networking components and analysis, namely for switched networks such as AFDX. Also a number of communication protocols and tools, developed for improving predictability and adaptivity in (industrial) networked embedded systems.

The cluster teams have been involved in many European projects, had strong interaction with industry and disseminated their work through active participation in world class conferences, workshops and schools.

1.3.4 Thematic Cluster: Hardware Platform and MPSoC Design

Year 1

- Platform and MPSoC Design

  The work has included:

  - Study of system design methodologies handling the dynamic nature of embedded systems and allowing predictability and optimal use of resources.
    Bologna, with ETHZ, has studied optimalisation-centric MPSoC design techniques.

The main goal of this work was to establish a common understanding of the MPARM framework developed in Bologna, and the DOL framework developed at ETHZ.
Bologna and Linkoping have studied a temperature power-optimization system. A temperature-aware dynamic voltage selection technique has been developed for energy minimisation.
Other work on design optimization for fault-tolerant distributed embedded systems is developed by Linkoping and DTU.

- Bologna and ETHZ have improved the design of a scavenger prototype, to perform automatic maximum power point tracking. They developed a compact model for small solar modules that accurately describes their behaviour over a wide range of irradiance conditions. Furthermore, they improved the efficiency of the DC-CD converter at the solar harvester.
- We (DTU) have studied programming models for MPSoC architectures as well as investigated the hardware/software interface between the processing elements and the interconnect network. We also have studied a component-based service model for early design space exploration and performance estimation.

- Platform and MPSoC Analysis

Work on analysis complete the design techniques above, with simulation, and performance analysis techniques:

- We studied techniques for performance estimation of distributed real-time systems, based on simulation, in particular for applications using heterogeneous task scheduling policies. We also studied performance analysis techniques for a MPSoC in collaboration with ST Microelectronics.
- An important work direction is modelling and performance analysis for multi-processor and/or networked systems.
- We studied relations between simulation-based and analytical methods for performance evaluation of distributed real-time systems. Based on experimental simulation results, we were able to draw interesting conclusions regarding the pessimism of formal approaches. The experiments were performed on FlexRay and CAN-based distributed systems.
- We also studied interesting relations between MPA (Modular Performance Analysis) and Timed Automata.
- We have extended the fault-tolerant process model, to consider a combination of hardware and software fault-tolerant techniques. We have proposed a method for computing the reliability of a system, taking into account: a) hardening levels in hardware; b) the re-execution levels in software; c) scheduling for sharing recovery slacks.
- We have worked on modelling and optimisation of a miniaturized solar energy harvester. We focused on the optimisation of two important metrics:
a) maximisation of the energy harvesting efficiency and b) the minimisation of the energy used for ineffective operations. A hierarchical control solution has been designed which overcomes several drawbacks of previously proposed approaches. A novel algorithm for approximate multi-parametric linear programming has also been proposed.

- We studied scheduling-based energy optimisation techniques for energy-scavenging wireless sensor networks.

Year 2

The principal Y2 results obtained are:

- Platform and MPSoC Design
  - MPSoC design and programming, including the development of:
    a) modelling concepts, methods and tools allowing cost-efficient mapping of applications,
    b) run-time resource management techniques,
    c) architectures on chip communication, in future many-core processors,
    d) MPSoC mapping tools, for multi-media and wireless applications,
    e) MPSoC architecture exploration.
  - The design of fault-tolerant distributed embedded systems,
  - Resource-aware, system-level optimisation, including,
    a) power optimisation,
    b) performance optimisation in energy harvesting systems,
    c) optimisation-centric MPSoC design.
- Platform and MPSoC Analysis
  - Modelling and analysis techniques for performance evaluation, including,
    a) performance estimation of distributed real-time systems, for control applications,
    b) unifying approaches for hierarchical scheduling,
    c) modelling of shared resources in multi-processor systems,
    d) modelling and analysis of adaptive systems,
    e) contract-based architecture dimensioning.
  - Analytic real-time analysis and timed automata,
  - Modelling and analysis of fault-tolerant distributed systems,
  - Verification of design properties of hardware architectures.

Year 3

The activities on Platform and MPSoC Design and Platform and MPSoC Analysis have been further integrated. The following is a list of some of the major achievements for year 3 showing collaboration between teams of the Cluster and other teams of the NoE. Details and more information can be found in the two activity reports by the cluster, one on design and one on analysis.

- EPFL-UNIBO: Interaction between EPFL and UNIBO was very active in Year 3. Major problems tackled include: 1) Network on Chips models and tools; 2) 3D integration models and analysis tools; 3) Study of NoCs for 3D integration. Exchanges with University of Bologna (UNIBO) continued from the previos years. Prof. Benini spent 2 months at EPFL as Visiting Professor.
- UNIBO-SSSA: UNIBO continued to collaborate with the Scuola Superiore Sant'Anna (SSSA) on predictability and modularity of MPSoC for Real-Time applications. The interaction has been realized by the integrating tools deveoped by the partners: the Erika RTOS from SSSA has been extensively used on the MPARM platform developed by UNIBO.

- ETHZ-TUBS: There has been intense cooperation about the coupling of two performance analysis methods, namely Symta/S and MPA. The corresponding tools have been connected and joint works on hierarchical event streams have been published.
- ETHZ-VERIMAG: The system for mapping algorithms onto MPSoC platforms (DOL) has been connected to the BIP system of Verimag with the advantage of a provably correct design flwo as well as a fast performance evaluation method that supports design space exploration.
- ETHZ-UNIBO: Based on a successful cooperation in terms of energy harvesting sensor networks, several further joint investigations of application control and hardware implementation have been performed, related to the area of CPS. Major extensions have been done in terms of harvesting in distributed settings and reward-based optimization strategies. These activities resulted now in a journal publication.
- TU Braunschweig-ETHZ: Interaction between TU Braunschweig and ETHZ has been in the area of performance analysis for multiprocessor systems with shared resources. Opportunities for improvement of modelling and analysis approaches were identified.
- KTH-ETHZ: Zhonghai Lu from KTH has visited ETHZ in the period November 2009-January 2010. The visit focused on performance analysis of embedded systems, on-chip communication and wireless sensor networks and has increased the mutual understanding of the two groups research efforts in this area. It generated several ideas for focused joint research topics.
- DTU-KTH: The two groups have a tight cooperation on the topic of system level modelling, which is also part of the SYSMODEL Artemis project. The KTH modelling framework ForSyDe is used as common basis for further developing system level modelling techniques. DTU has develop the discrete time modelling domain of ForSyDe further, allowing for faster simulation as well as for parallelization of the simulation kernel, which is a key element when modelling wireless sensor networks (CPS). KTH has focused on synchronous, the untimed and continuous time domain. Mikkel Koefoed Jakobsen from DTU has visited KTH twice during 2010 for periods of several weeks to a few months to foster the joint work on the ForSyDe framework. One of the new demonstrators for this work, is a medical audio-device for adjusting hearing aids.
- CEA LIST and UNIBO have continued their collaboration for the definition and design of a Software Runtime Architecture for the management of many-core components. They also work on the design of efficient hardware support for the execution of this runtime software. This runtime SW is distributed to other prtner within the framework of SMECY project. A joint publication has been submitted on part of this collaboration.
- LINKÖPING-DTU: Linköping and DTU have continued their work on fault tolerant embedded systems. This has resulted in joint development and publications. Prof. Paul Pop from DTU has visited Linköping.
- LINKÖPING-LUND: Linköping has a close cooperation with Lund (Artist design partner Cluster: Operating Systems and Networks) in the area of modelling and QoS optimisation of control applications. This has resulted in joint development and publications. Soheil Samii and Anton cervin have visited Lund and Linköping, respectively.
- IMEC-NTUA: IMEC and NTUA have been collaborating on several MPSoC topics incluiding a framework for automatic parallelization, static and dynamic memory optimization in MPSoC platforms, runtime system exploration for multi-standard Wireless MPSoC.

- IMEC-KTH: IMEC and KTH continued their collaboration in the context of the European project MOSART (http://www.mosart-project.org). The Co-Ware virtual multi-core platform developed by IMEC was transferred to KTH for integration of the NoC architecture in the platform model.
- IMEC-NTNU: there has been cooperation on data value driven scenario identification and reuse of epilepsy detection kernel as additional biomedical demonstrator for scenario related research.

Year 4

The Cluster has continued its efforts to establish an integrated modelling and design methodology that can take into account predictability and resource-awareness with focus on efficiency. This work has benefited from fruitful collaboration with the Cluster on Modelling and Validation and Timing Analysis as well as from the transversal activities on design from adaptivity and predictability.

Main results can be summarized as follows:

- Fault tolerant distributed embedded systems: We have developed results for handling both processor and communication faults in distributed real-time systems for automotive applications, based on CAN or FlexRay communication.
- Performance analysis methods: TU Braunschweig and ETH Zurich have developed very original and relevant results. They have collaborated to establish a method for coupling the tools SymTA/S and MPA. Relying on different analysis techniques each of the two tools can be individually used to evaluate the performance of embedded real-time systems. The interface developed for tool coupling now allows combining the strengths of the two tools. Evaluations have been jointly performed and the work resulted in joint publications.
- MPSoC design: Major activities on MPSoC design have focused on application parallelization, platform mapping, memory hierarchy management, application scenario exploitation, and run-time resource management, including reconfigurable systems. The outcome of these 4 years was the development of related tools, tool integration in tool chains in collaboration with several ArtistDesign partners, and highly referenced publications.
- Energy harvesting: We have developed new node level scheduling techniques (UNIBO and ETHZ) as well as network level routing algorithms (DTU), and have demonstrated that these techniques can lead to considerable extensions of the lifetime of the network. One specific outcome is the founding of the company WISPES srl (Wireless Self-Powered Electronic Systems) that aims at providing technologies and devices able to add wireless communication and local computation to the customer's monitoring and sensing activities.
- Temperature and energy aware optimization: EPFL has developed a novel online thermal management policy based on dynamic voltage and frequency scaling for high-performance 3-D systems with liquid cooling. The approach is able to gain up to 50% as compared to current state-of-the-art thermal control techniques.

Finally, the Cluster has an impressive record of joint publications, invited talks, analysis and design tools and industrial collaborations.

On the widespread use of multi/core and many core processor architectures

a. The Need for Multiprocessing

Embedded devices perform an increasingly wide variety of high and low-level tasks. With the growing availability of high speed mobile and Wi-Fi networks, embedded devices will increasingly be used for performance-intensive tasks that were previously handled by traditional non-embedded computers. For example, the next generation of smartphones (called “Super phones”) and tablets will be used for a wide variety of tasks such as playback of high definition 1080p videos, Adobe® Flash®-based online gaming, Flash-based streaming high definition videos, visually rich gaming, video editing, simultaneous HD video downloads, encode and uploads, and real-time HD video conferencing.

Even for high-end devices, the current generation of embedded processors is not designed to deal with this tidal wave of high performance use cases. The quality of performance for devices based on single core CPUs rapidly degrades when users run several applications concurrently, or run performance intensive applications. To improve CPU performance, engineers employ several techniques, such as using faster and smaller semiconductor processes, increasing core operating frequency and voltage, using larger cores, and using larger on-die caches.

Increasing the size of the CPU core or cache delivers performance increases only up to a certain level, beyond which thermal and heat dissipation issues make any further increase in core and cache size impractical. From basic semiconductor physics we know that increasing operating frequency and voltage can exponentially increase power consumption of semiconductor devices. Even though engineers may be able to squeeze out higher performance by increasing frequency and voltage, the performance increase would drastically reduce battery life. In addition, processors that consume higher power would require larger cooling solutions resulting in an undesired expansion in device size. Therefore, increasing the operating frequency of the processor to meet the ever-increasing performance requirements of mobile applications is not a viable solution for the long run.

b. Trends

For several decades processors were primarily single core architectures and CPU makers increased performance by increasing operating frequencies, core sizes, and using smaller manufacturing processes allowing more transistors in the same chip area. However, CPU manufacturers realized that the continual increase in frequency and core sizes caused exponential increases in power consumption and excessive heat dissipation. Therefore, CPU makers developed multi-core CPU architectures to continue delivering higher performance processors, while limiting the power consumption of these processors. Most desktop and notebook systems today use either a dual or quad core processor and consume significantly lower power than their single core predecessors. But mobile devices like smartphones and tablets benefit even more from multi-core architectures because the battery life benefits are so substantial. Dual-core processors and quad-core are now the standard.

Embedded application processors are facing the same performance and power challenges that desktop and notebook CPUs faced a few years ago. Applications are already stretching the capabilities of current single core mobile processors. To further increase the performance and stay within embedded power budgets, it is likely that most embedded processors will eventually have multi-core processors. Mobile operating systems such as Android, Windows®CE, and Symbian are capable of operating in a multi-core environment, and have the features required to efficiently harness the multiple processing cores of the underlying hardware.

1.3.5 Transversal Integration activity: Design for Adaptivity

Year 1

We have worked mainly in two complementary directions: a) Study of architectures and algorithms for ensuring adaptivity; b) Study of modelling and analysis techniques for adaptive systems:

- We studied a symbolic quality control technique for multi-media applications. Adaptivity is ensured by using a controller, which moniotors system execution and adapts quality parameters of its functions so as to meet hard real-time contraints.
- We studied adaptive energy management techniques in clusters of wireless sensor nodes. They allow tuning the application parameters according to the time-varying amount of harvested energy.
- We studied a reference architecture for self-configuring embedded systems in collaboration with Volvo. Algorithms suitable for runtime configuration management, load balancing, and quality of service have been developed and adapted to automotive applications.
- We designed adaptive techniques to enhance real-time support of IEEE 802.11.e networks. For such networks, we developed protocols to enhance the resilience to interference, and to provide an estimation of a relative localisation based on the radio frequency signal.
- We studied techniques allowing the design and performance analysis for multi-mode systems. We also studied online performance analysis techniques for distributed systems. A novel distributed algorithm for control of the global analysis flow has been proposed.
- Several partners have collaborated in the STREP projects FRESCOR and ACTORS to develop an infrastructure for adaptive scheduling of real-time applications.

Year 2

The principal Y2 results obtained are:

- Analysis techniques for adaptive systems, including
  a) dynamic changes in real-time parameters,
  b) assignment of real-time parameters to control tasks,
  c) timing analysis and sampling mechanisms, for event-driven control systems.
- Adaptive design techniques, including:
  a) algorithms for QoS-aware, cooperative systems,
  b) adaptive topology management, to combine energy efficiency and QoS,
  c) in-system self-optimization for real-time systems,
  d) reconfigurable self-organizing and self-healing hardware platforms,
  e) adaptive energy management.
- Modeling and analysis of adaptive systems,
- Support for Adaptivity in distributed systems.

Year 3

The partners have organized several workshops and meetings, including WARM 2010. The meetings act as the interface between the different clusters on issues related to embedded system adaptivity.

The partners have contributed to education about adaptive and feedback-based approaches. There are also several contacts between industry and academia within the activity, e.g., collaborations involving NXP, Ericsson, Volvo, IMEC, and Evidence just to name a few.

A major challenge for this activity continuous to be how to integrate the more hardware-oriented partners from, e.g., the MPSoC cluster with the more software-oriented partners from the OS and networks cluster. Currently the activity is dominated by partners from the latter cluster.

The members of the activity are organizing a special issue on Adaptive Embedded Systems for Real-Time Systems Journal with Årzén (ULUND) as guest editor. The deadline for submissions is Sep 2011 which fits quite nicely with the end of ArtistDesign, making it possible for the members of the activity to submit their work there.

We provide a list of technical achievements of the partners, both jointly and individually, during Year 3, structured in three groups: adaptive resource scheduling, adaptive networking, and hardware adaptation. In the first and largest group we also include work on modelling and analysis relevant to adaptation.

- Adaptive Resource Scheduling
  - Adaptive and feedback-based resource management (SSSA, ULUND, TUKL, Evidence, Ericsson)
  - Adaptive resource management for uncertain execution platforms (ULUND, Ericsson)
  - Feedback control of computing systems (ULUND)
  - Theory of distributed performance analysis (TU Braunschweig)
  - In-system sensitivity analysis for real-time systems (TU Braunschweig)
  - Change impact analysis (UYork)
  - Parametric WCET analysis (MDH)
  - Runtime management of cache-related preemption delay (IPPorto)
  - Fault tolerance in adaptive cooperative systems (IPPorto)
  - Dynamic behavior of embedded systems (IMEC, NTUA)
  - Adaptive control of MPEG-4 decoding (TUKL, ULUND)
  - Improving real-time BIP (Verimag).
  - Adaptation in service-oriented architectures (UPM)
  - Adaptive servers with guarantees (ETH Zurich, SSSA)
  - Adaptive power management (ETH Zurich, SSSA)
  - Sampling mechanisms for event-driven control systems (UPC, ULUND, SSSA)
  - Feedback scheduling vs. event-driven control (UPC)
  - Optimal online sampling period assignment (ULUND, UPC)
- Adaptive Networking
  - Adaptivity in wireless networks (UPorto, UCatania)
  - Adaptivity in distributed systems (UPorto, MDH, UAveiro, UPC)
  - Adaptive management in energy harvesting systems (ETH Zurich, UBologna)

- Adaptive energy management of wireless smart camera networks (UBologna)
- Adaptive TDMA bus allocation and elastic scheduling (UBologna, SSSA)
- Fault Tolerant and Reliable Communication Platforms (KTH)

- Hardware-Based Adaptivity
  - eDNA: Reconfigurable self-organising and self-healing hardware platform (DTU)
  - Adaptive allocation of applications on MPSoC platforms (ETH Zurich, SSSA)

Year 4

The work done includes numerous highlights:

- Scheduling analysis: Efficient and effective scheduling analysis for fixed priority systems has been developed that takes into account tasks arriving and leaving the system. Furthermore, a new method for allocation and scheduling of parallel tasks in soft-real time systems (multimedia decoding) in the presence of post-silicon, process and ageing induced variability in a nominally homogeneous target multi-core platform has been developed.

- Memory: Dynamically adaptable memory architectures for supporting dynamic real-time process loads have been developed.

- Collaboration frameworks: An adaptable cooperation-based framework for networked embedded systems with heterogeneous nodes has been developed, allowing constrained devices to cooperate with more powerful (or less congested) neighbours, to meet allocation requests and handle stringent constraints, opportunistically taking advantage of global resources and processing power.

- Service adaptation: Techniques have been developed for adapting the service request handling behaviour to the specific requirements of the services in Service Oriented Architectures (SOA). CPU contracts are used to ensure sufficient computation time for dealing with services with special requirements.

- Run-time resource management: An adaptive resource manager for distributed embedded systems aimed at multimedia applications, e.g., broadcast management systems, was developed. Considerable savings in power consumption, hardware cost and system size were reported in an industrial case study. Parallel to this a QoS based adaptive resource management system for homogeneous multicore platforms was developed.

- Run-time analysis: A distributed approach for in-system run-time performance analysis of embedded systems, complemented by a framework enabling access control and runtime-optimization through the use of distributed algorithms.

- Sensor networks: New approaches to adaptive energy management of energy harvesting system using solar cells have been developed. Based on a prediction of the future available energy, the application parameters are adapted in order to maximize the utility in a long-term perspective.

- Control techniques: A new method for optimizing the timing parameters of real-time control tasks in resource-constrained embedded systems has been derived. Also, new

feedback scheduling techniques and new event-driven sampling mechanisms have been proposed.

- Adaptivity in networks: Here various ways of adapting a communication channel to varying application requirements or environmental conditions to enhance the efficiency of medium utilization have been proposed. For controlled access networks with isolated virtual channels the guaranteed bandwidth and latency can be adapted online using the Flexible Time-Triggered (FTT) paradigm on switched Ethernet, either with COTS switches (FTT-SE protocol) or enhanced ones (FTT-enabled switch).

- Programmable hardware: A new type of ultra-fault-tolerant FPGA named the eDNA architecture has been conceived all the way from development of the concept, to the implementation of a prototype, to test in a space related case study NASA JPL.

- WCET analysis: Parametric WCET bounds, where the WCET bound depends on the values of certain inputs, can be used in adaptive real-time systems where the scheduling of tasks adapts to external factors such as varying data sizes affecting the running times of tasks. A general method for parametric WCET analysis, which combines a number of advanced symbolic techniques including relational abstract interpretation, counting of integer points in polyhedra, and parametric integer programming has been developed and implemented in the WCET analysis tool SWEET.

- Reference architectures: A reference architecture for automotive embedded systems that addresses the needs for flexible and automatic run-time reconfiguration has been proposed. The research focus was the development of technical support in terms of middleware services for a closed adaptation of distributed embedded systems. In addition to the reference architecture an information model of the control parameters that represent the target system configuration alternatives, environmental parameters, and internal conditions has been defined and a functional design has been performed.

Research Trends and Vision on “Design for Complex Adaptive Systems”

Real-time systems constitute a notable share of today’s embedded computers that needs special attention. The design of robust and fault-tolerant real-time systems is a highly active research area that has produced numerous approaches for evaluating and increasing system robustness against selected fault scenarios. These methodologies can be applied throughout the design process of an embedded system and yield systems that are highly robust against a selected set of disturbances in the field. Future embedded systems, however, will undergo an evolution in both hard- and software configuration during their lifetime. In the automotive industry, it is already common to update or add software components during the lifetime of a product, producing a variety of software configurations in the field. To ensure functional and temporal correctness of all possible configurations, OEMs have to maintain a complex versioning database and perform exhaustive testing to cover the whole configuration landscape. This already constitutes a problem today, which will grow into a major challenge in the future.

Designing embedded systems to be robust and fault-tolerant will not ultimately solve this problem, as the evolution an embedded system goes through during its lifetime cannot be foreseen at design time. Hence, embedded systems need to be adaptive to changing conditions, in the sense that they need to be able to meet given requirements including safety, security, and performance, in the presence of uncertainty in its external environment or internal execution platform. Adaptivity can be seen as a means for enforcing predictability in the presence of uncertainty.

The uncertainty can be viewed as the difference between the average and the worst-case behavior of a system and its environment. The trend in embedded system is towards drastically increasing uncertainty due to, e.g., execution platforms with increasingly sophisticated HW/SW architectures (layering, caches, multiple cores, speculative execution etc), increased connectivity with complex and non-deterministic external environments, increased amount of difficult-to analyze software, and increased variability with respect to use cases.

One technique for achieving adaptivity in particular in software-based systems is feedback. In many embedded systems worst-case designs are unfeasible for several reasons. One of these is the over-provisioning of resources that this typically implies. Other reasons are uncertainties associated with worst-case resource utilization estimates and on-line changes in objectives, external conditions and use cases. In a feedback-based resource management system, the allocation of resources is based on a comparison of the actual resource utilization by, e.g., a set of activities or tasks, with the desired resource utilization. The difference is then used for deciding how the resources should be allocated to the different activities. The decision mechanism constitutes the feedback controller in the scheduling scheme. Feedback control makes it possible to deal with uncertainties and variations in a controlled way.

Feedback scheduling is primarily suited for soft real-time applications and adaptive real-time applications, where missing one or more deadlines does not jeopardize correct system behavior, but only causes performance degradation. For this type of systems, the goal is typically to meet some Quality of Service requirement. The adaptive class of real-time systems is a suitable description for a many practical applications. This includes different types of multimedia applications, but also many control and signal processing applications. An important research trend here is how to best model embedded computing system from a control perspective. Different model formalisms can be considered, from pure discrete event based models to fluid continuous-time approximative models.

The research trends related to adaptivity in embedded systems are numerous since adaptivity is crosscutting. In hardware-based oriented embedded systems there are work performed on modeling and hardware generation for adaptive processes and applications. Emerging architectures such as partially reconfigurable, either fine-grained or coarse-grained, FPGAs provide a huge potential for adaptivity in the area of embedded systems. Since many system functions are only executed at particular points of time they can share an adaptive component with other system functions, which can significantly reduce the design costs. However, adaptivity adds another dimension of complexity into system design since the system behavior changes during the course of adaptation. This imposes additional requirements on the design process, in particular system verification.

In the software-oriented part of embedded systems there is also a considerable work on computational models that allow for adaptivity, how adaptivity can be provided in component-based architectures, adaptive task models for scheduling, program language constructs supporting adaptivity, and run-time support for adaptive resource management from operating systems, middleware, and communication networks. The resources in the latter case typically include clock cycles, memory, communication bandwidth, and energy, but could in general also include other resources which are allocated dynamically.

The cyber-physical system (CPS) trend has also strong relations to the work within this activity. One example of this is holistic design approaches that integrate physical aspects such as power and heat, with cyber aspects.

1.3.6 Transversal Integration activity: Design for Predictability and Performance

Year 1

The technical work on Predictability has intersected work in all the Thematic Clusters.

- Modeling and Validation of component –based systems
  - We studied the concept of predictability in relation with robustness, and identified two major challenges in embedded systems design. These challenges require a rethinking of the conventional, purely discrete and purely functional foundation of computing.
  - We worked on modular performance analysis techniques, based on real-time calculus for systems with cyclic dependencies. We integrated a contract-based scheduling framework with a component-based technology.
- Timing Analysis and Compiler Techniques
  - The main contribution is work on relations between Timing Analysis and Timing Predictability. A definition of predictability for cache architectures has been proposed, and the relative competitiveness of 4 different cache replacement strategies has been analysed.
  - We also investigated WCET analysis techniques for cooperative task scheduling. A method guiding developers of an embedded system to select optimal pre-emption points is under development.
  - We also studied parameteric timing analysis techniques that overcome usual limitations of analysis techniques requiring the maximum number of loop iterations to be known statically.
- OS/MW/Networks
  - Our work addressed various issues, including: Integrating scheduling analysis and model checking; influence of abstractions on the schedulability analysis of distributed real-time systems, time-predictable operating systems.
- Architecture and System Design
  - We studied techniques allowed predictability for fault-tolerant embedded systems and predictability for multi-processor MPSoC architectures.
  - We also developed a set of parametyerisable models of L2 Caches and integrated them in an accurate virtual platform environment.
  - Finally, we are designing a Precision Timed (PRET) architecture based on a reactive processor, coupled with a MicroBlaze general purpose processor.

Year 2

The technical work on Predictability has intersected work in all the Thematic Clusters.

- Modeling and Validation of component –based systems
  - We studied the concept of predictability in relation with robustness, and this work is a continuation of work in Year 1. We identified two major challenges in embedded

systems design. These challenges require a rethinking of the conventional, purely discrete and purely functional foundation of computing.

- We worked on methodologies for designing component-based systems, and participated in the standardisation of the MARTE-UML profile.

- Timing Analysis and Compiler Techniques
  - The main contribution is work on relations between Timing Analysis and Timing Predictability. We identified heuristics to derive operating mode candidates from source code and a procedure to exploit mode information to arrive at mode-specific WCET bounds. We also started to explore strategies for eliminating timing anomalies by different code generation techniques.
  - We investigated WCET analysis techniques in the presence of context switches.
  - Finally, we integrated Timing Analysis and Compilation techniques for optimizing code generation.

- OS/MW/Networks
  - Integrating scheduling analysis techniques and model checking.
  - Partitioning the shared caches on multi-core processors for timing predictability. We have developed a sufficient schedulability test for non-preemptive fixed priority scheduling for multi-cores, with shared L2 cache, encoded as a linear programming problem. Furthermore, we have compared several scheduling models for multi-processor systems.

- Architecture and System Design
  - We studied techniques allowing predictability for fault-tolerant embedded systems and predictability for multi-processor MPSoC architectures. In particular, we developed power prediction algorithms for an energy harvester.
  - We also developed a new bus model for MPSoC system bus analysis and optimisation.
  - Finally, we developed the Precision Timed (PRET) language, which allows determinism and time predictability.

Year 3

The technical work on Predictability has intersected work in all the Thematic Clusters. We give some examples of resulting progress on several topics:

- Novel collaborations within the PREDATOR project 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.
- 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.

- Work on reconciling timing analysis with compilation includes the development of the WCET-aware compiler WCC by TU Dortmund, in collaboration with 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. WCC has been extended towards code generation and optimization for multi-process systems

Finally, we organized a workshop on predictability and performance at DATE 2011.

Year 4

The Predator project has made strong progress in its attempt to reconcile Predictability with Performance. The integration of the AbsInt timing-analysis tool aiT with the WCET-aware compiler of TUDortmund is described separately. Another recent achievement of the project concerns the determination of context-switch costs, which provides support for schedulability analysis for preemptive scheduling strategies. Insights into the predictability properties of architectural features have found their way into the embedded-systems industry, e.g., as a result of collaboration in European projects. These insights, however, are still at odds with trends at the processor manufacturers side. Suppliers of time-critical embedded systems cannot find platforms with the required predictability properties on the market.

The trends to multi-core platforms presents a significant challenge to the building of predictable and performant systems, and there is still significant hesitation to migrating embedded systems to multi-cores. Significant advances on isolation and analysis techniques have been made (to a large extent by ARTIST-Design partners): progress is made, e.g., in the area of deterministic access protocols and controllers for shared resources such as buses or memory. However, the worst-case delay used in safe approximations is still often too high to be acceptable.

A good collection of insights was gathered at the PPES workshop, organized by ARTIST-Design, jointly with Predator and Merasa, as a satellite event of DATE 2011 in Grenoble. Overviews about architecture and software issues were given, e.,g including a survey on predictability and performance requirements in avionics systems, and a template for, partly analytically, partly intuitively, estimating the predictability of hardware features was presented.

During year 4, development of support for the MARTE standard (initiated during ARTIST2), led by U. Cantabria, has provided increased support for scheduling and code generation. The work on integration between timing analysis tools has matured: several of the leading timing analysis tools have been integrated by efforts in the All-Times project (described in the report on Timing analysis).

A notable trend during Year 4 has been the work on reconciling predictability with performance, developing techniques for optimizing performance along several dimensions (e.g., combing WCET with average-case timing). Work in this direction (by Bologna, ETHZ, Linköping, Trento) has considered different forms of multi-objective optimization of embedded software; such possibilities also exist in the WCC compiler. Another increasingly important topic has been to make scheduling and timing analysis robust to inaccuracies in assumptions about, e.g., execution times, interferences, etc.

Work on the integration of timing analysis and compilation, in the context of the WCC compiler, aimed at removing some of the earlier restrictions. The work started at TU Dortmund considering WCET-aware basic block reordering has been finished. Unconditional branches are avoided and the prediction of conditional branches is supported by the developed techniques. A genetic approach applies evolutionary algorithms considering the WCET of the program to optimize as fitness value with the costs of high optimization times. Thus, an integer-linear programming-based approach has been developed which determines the optimal order of basic blocks and also takes the branch prediction into account [PKFM11]. Furthermore,

WCET-aware cache locking and code positioning has been improved. The integration and enhancement of a framework for the static analysis of software and hardware as announced in last year's report has been advanced. The extension of WCC's native analysis capabilities allows for novel approaches especially in the domain of multitask- and multicore-aware compilation. A much higher degree of control over system states directly affected by optimization decisions can thus be achieved. The primary effort was made in the direction of tightening timing estimations and the evaluation and improvement of cache analysis techniques with a focus on improving compiler optimizations [KFM11].

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

Research Trends and Vision on "Design for Predictability and Performance"

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 are being investigated. A natual approach is to enable a priori predictability analysis and to support mappings to platform architectures that preserve predictability. Investigations of how modeling and analysis techniques extend to non-traditional system structures, including parallel, distributed and networked architectures, for which predictability is more difficult to achieve is currently a very active research area. A more fundamental research direction is to develop precise definitions and characterizations of central concepts, including predictability and robustness. Another overarching research theme is to explore 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 issue 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. They introduce new opportunities for parallelism and communication, in order to enable higher performance. At the same time they pose new significant challenges for the building of predictable systems. One major problem is that multiprocessor techniques typically introduce new forms of interference, even between activities that are logically unrelated. For instance, execution on one processor may interfere with the timing of an unrelated activity on another processor through a shared L2 or L3 cache. The timing of memory accesses can be affected by unrelated memory accesses of other processors through interference on a shared bus, and so on. It is not yet clear how to build predictable and performant systems on multicore platforms.

1.3.7 Transversal Integration activity: Industrial Integration

Each of the ArtistDesign Thematic Clusters (WP3-WP6) is important per se for advancing the state-of-the-art in embedded system design. However, if we wish to have a strong impact on industry and society at large, the results of the thematic clusters have to be harmonized in an overall design flow that can sustain the industrial embedded design chain from conception of the product to its implementation.

The chains vary in length and players according to the industrial segment addressed: for example, the design chain in automotive electronics starts with the car maker (e.g., BMW, Daimler Chrysler, Peugeot, Fiat), goes through the Tier 1 suppliers (e.g., Contiteves, Bosch, Magneti Marelli) and connects to the Tier 2 suppliers (e.g., FreeScale, ST, Infineon, Hitachi). It often includes IP providers such as programmable cores, RTOS and software development tool providers and design service companies. In the mobile communication domain, the chain starts with the application developers (e.g., gaming and video content), includes the telecommunication operators (e.g., Telecom Italia and Telefonica), the device makers (e.g., Nokia and Ericsson), the silicon makers (e.g., TI, Qualcomm and ST) and outsourcing manufacturing companies (e.g., Flextronics).

Today, there is stress in the chain as the technology advances may create opportunities to redefine the roles of the various players. In addition, the system integrators are often faced with an almost impossible task of composing their design out of parts supplied by companies whose design methods and standards are widely different and about which they have limited or no

information. There is a need for an all-encompassing approach to system design that can make an entire industrial segment work as a virtual vertically integrated company. The benefits of these flows and methods are obvious as they provide shorter time to market and better quality designs but require a will of the industrial segment to work together towards this goal. In the automotive domain, Autosar is an excellent step in that direction. Other industrial segments are less cohesive in searching for a unified approach to design. In addition, society concerns such as energy, health and environment conservation, are offering new business opportunities for emerging technologies such as wireless sensor networks. The difficulty in these new opportunities resides in lack of standards and of experience with new communication concepts and, last but not least, in security.

We believe that all the thematic clusters bring something important to all industrial segments, but we need to pay attention to the way the results obtained by the clusters are formulated. Integration is a matter of modelling and providing interfaces that guarantee that the properties of the components are maintained after integration. Integration takes two forms: an horizontal one where different IPs coming from different companies or from different design groups in the same company have to be assembled; a vertical one, where the requirements are clearly and possibly formally communicated from a higher level player to a lower level one and where the information about the capabilities and limitations of the IPs are unambiguously communicated from the lower level to the higher level. The ultimate goal of this activity is to provide the “meta rules” according to which the design transformations are carried out and interfaces are built and hence to provide strong guidance to the clusters to make their results more relevant and applicable. Understanding the roles and dynamics of an existing, well-established, vertical industrial segment is a complex task. We could only imagine the complexity of industrial segments that are coming together in these years. While we do target some industrial domain to be the driver for this activity, we understand that our research is going to be more relevant and better quality if we can distil some common traits of these domains and work with those to choose at a later date which particular chains to address.

The transversal activity hence has two prongs:

- to dive into particular vertical industrial segments and package design methods out of the thematic cluster results for the segments;
- to identify some important common features among verticals and work towards developing methods to address these topics.

We note that the two concerns that are also part of the Transversal JPRAs (predictability and adaptability) are common to almost all industrial concerns: For this reason, they provide a framework to start the work on integration driven by industrial applications. Predictability has been a goal since the beginning of the modern industry: predicting the capabilities of existing components allows companies to come to market faster with new products and prevents taking dead ends; predicting the effort needed to develop parts of the design and their correct integration prevents early recalls and associated costs. The faster is the dynamics of the industry, the more important is to have predictability in design.

Adaptability is the property of a design to be adapted to changing environments and working conditions. Reconfigurability, programmability, dynamic restructuring are all facets of adaptability. Novel approaches to communication could benefit greatly from adaptability. In fact, much research is being carried out to design devices that could sense available bandwidth and adapt the communication protocol to the most convenient band at the time.

We believe that it will be eventually easier to compose the vertical design industrial flows once these two sub-flows have been examined and results obtained. In addition, being generic concerns they do not require effort from the academic partners to understand the modus operandi of entire industrial segments and offer a shorter time to results.

The vertical industrial segment motivated prong will begin by bringing up-to-speed the largest possible number of participants to the logic of the design chain by organizing workshops for discussion with the participants to the chain.

We proposed at the onset of the activity to target Automotive, Nomadic and Health Applications as potential vertical segments where we have a range of maturity from well-established (automotive) to emerging (health). At the 2008 meeting in Rome of the ArtistDesign partners, the three vertical markets of interest were identified as:

Automotive/avionics since we noted a strong similarity in the overarching issues faced by these two industrial segments that are driven by safety concerns and have to consider distributed implementations;

Health applications with particular emphasis on equipment design and manufacturing and a new thrust in the use of emebedded system design methodologies to synthetic biology;

Energy efficient buildings, a novel field of great interest to the European Community as well as to the rest of the world as 30% of energy consumption is considered to be in commercial buildings.

These applications address an established area of excellence of European Industry where international competition is fierce, an area of growth where again European Industry has a strong position but where the dynamics are fast and new applications are envisioned in strategic areas such as elderly care, and a new area with great potential where energy conservation concerns are going to place a great political emphasis. In addition, we believe that synthetic biology is going to have a fundamental role for the foreseeable future in the definition of new organisms to foster the creation of new drugs as well as new materials. Given the nature of this work, the main participants in the cluster are the groups that have industrial vocation such as ESI, OFFIS, and IMEC.

Year 1

This activity groups a set of industrial interactions and collaborations with ArtistDesign teams. The long-term goal is to understand industrial design methodologies and identify the research results that could be applied in these methodologies.

The work this year has consisted in organising a few high-profile meetings with industry (eg: Embedded Systems: Industrial Applications ’08) as well as joint workshops and technical meetings.

At this point, these constitute a rich set of events and interactions, which need to be structured and which need a more specific focus.

Year 2

This activity groups a set of industrial interactions and collaborations with ArtistDesign teams. The long-term goal is to understand industrial design methodologies and identify the research results that could be applied in these methodologies.

The work this year has consisted in :

- organising a few high-profile meetings with industry (eg: SEEC ’09, WESH ‘09,) as well as joint workshops and technical meetings. SEEC in particular was extremely successful, with a high international visibility far beyond European borders.
- a wide array of collaborations with industrial partners and ArtistDesign partners, including general frameworks for system-level design, automotive applications, applications for chip design, smart energy-efficient buildings, and wireless communication.

Year 3

This activity groups a set of industrial interactions and collaborations with ArtistDesign teams. The long-term goal is to understand industrial design methodologies and identify the research results that could be applied in these methodologies.

The activities include both technical achievements and dissemination work on the follwowing: General Frameworks for system-level design; Applications to the Automotive Sector; Applications to Chip Design; Applications to Buildings; Applications to Wireless communication technology; Timing Analysis and Predictability; Other Applications.

The level of energy at the meetings organized to foster industrial integration was excellent. In 2009, we proposed the change from Nomadic to Energy Efficient Building has had a re-sounding success. This theme is of increased interest to the European community in response to energy conservation concerns. In this respect, in 2009 a detailed plan was drafted for meetings to be held in 2010 and a modus operandi that included international interaction. The GREEMBED Conference was a result of these efforts. In 2010, we launched a new direction in the area of Synthetic Biology, with the sponsorship and participation to the 2010 International Workshop on Bio-Desing Automation. This area is bound to have a strategic impact on research world-wide. The meetings were very well attended and strong positive feedback was received also from some of the companies involved.

Year 4

The level of energy at the meetings organized to foster industrial integration was excellent. This theme is of increased interest to the European community in response to energy conservation concerns.

Cooperations with institutions:

- ESI and Healthcare - technical support for health care at home
  safety, performance, reliability, evolvability of large medical equipment, such as X-ray and MRI systems. Projects include:
  - ALWen – with several SMEs and research institutes
    Ambient Living with Embedded Networks: body sensors, ambient sensors, wireless networks, and telemedicine for health care
  - Care4ME – ITEA
    advanced medical imaging and decision support methods
  - Allegio – Philips Healthcare, Axini, universities
    safety, performance, and reliability of X-ray equipment
  - Darwin – Philips Healthcare,universities
    evolvability of medical equipment
- KTH and ICES (Innovative Centre for Embedded Systems)
  System architecture, Softvare verification, Methodology
  - A KTH-based centre for embedded systems
  - A proactive effort in the embedded systems area
  - ICES acts as a Network & Catalyst
  - 6,4 times gain factor
  - Current members
  - ABB, Enea, Ericsson, Freescale, Mathworks, Prevas, Scania, Semcon/EIS, Stoneridge, ÅF

- Four KTH schools and several groups encompassing
- Computer science, SW engineering, Computer engineering, Automatic Control, Communication, HMI, Mechatronics

Cooperations with large EU projects:

CESAR

Cost-efficient methods and processes for safety relevant embedded systems
55 partners, industry driven (78% Industry / 22% Academics)

- Objective: Interoperability Standard (IOS) for the development of relevant embedded systems, which are compliant to domain-specific safety standards
- Approach: Common semantical basis(Meta Modeling)
- Focussing on: Requirements Engineering, Component-based design, Development process
- Implemented as a Reference Technology Platform (RTP)
- Innovative analysis techniques based on the IOS
- Cross domains, but with domain specific tailoring

MBAT

Combined Model-based Analysis and Testing of Embedded Systems

- Objectives: Provide Europe with a new leading-edge Reference Technology Platform for effective and cost-reducing Validation and Verification of Embedded Systems
- Approach:
  - Based on meta models and compatible components to enable construction of customized System Analysis & Test Environments
  - Combined Model-based Analysis & Test Methodology including innovative analysis and test case generation techniques on different development levels
- Tool support based on an interoparability standard (RTP)
- Compliant to CESAR RTP
- Industry driven (cross domains)
- Business needs
- Use case and derived requirements

DANSE

Designing for Adaptability and evolutioN in System of systems Engineering (DANSE)

- Focus
  - Smarter and adaptive power grids
  - Reduction of buildings overall energy consumption
  - Optimization of industrial processes with timely information about energy cost/availability.
- Enabling technologies

- Distributed and pervasive sensing and monitoring
- Smart and effective communication solutions
- Predictive control of the environments
- Innovative way to generate, store and distribute energy

- Participation
  - 6 lecturers
  - 21 students (post-docs, PhDs, and young researchers from industry)

Other Advanced Topics and Opportunities

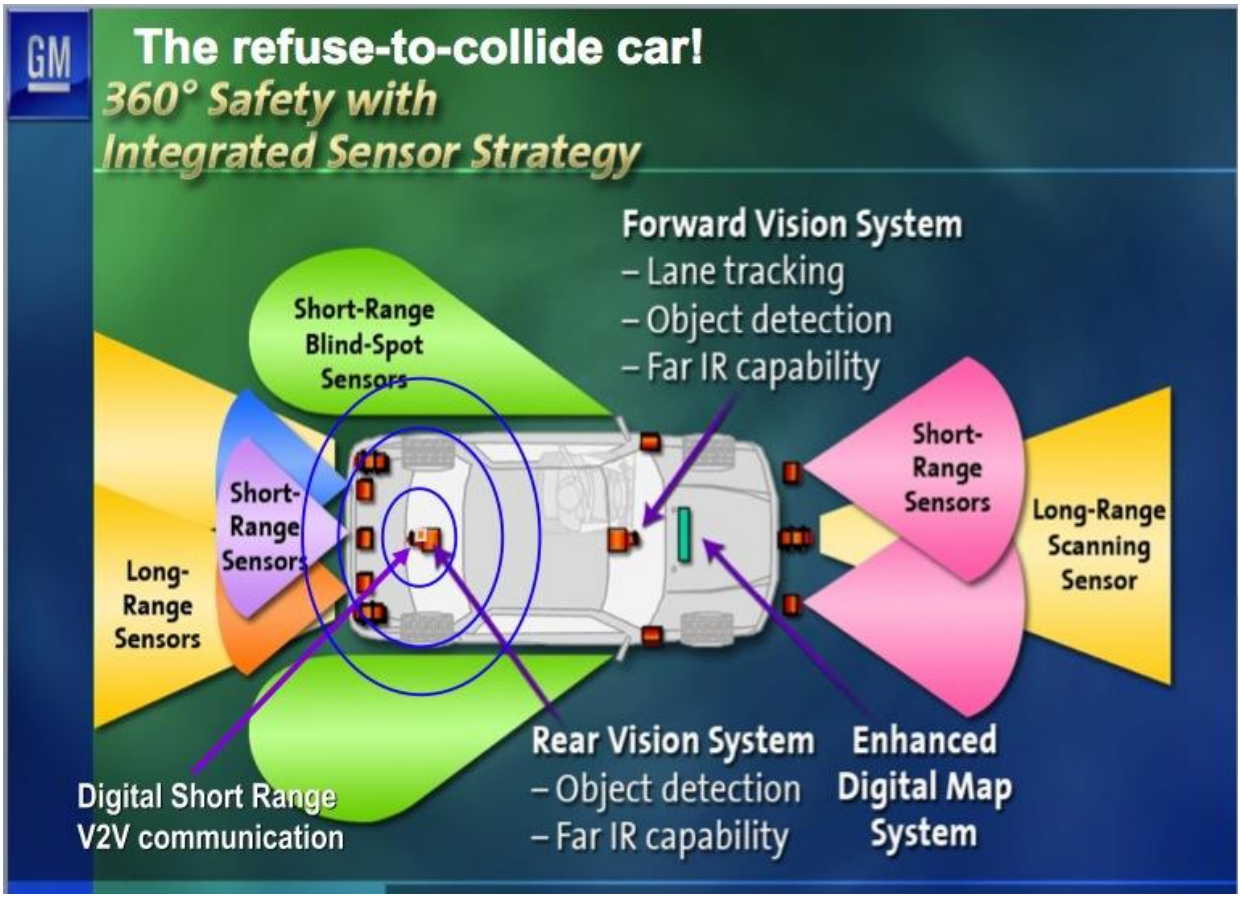

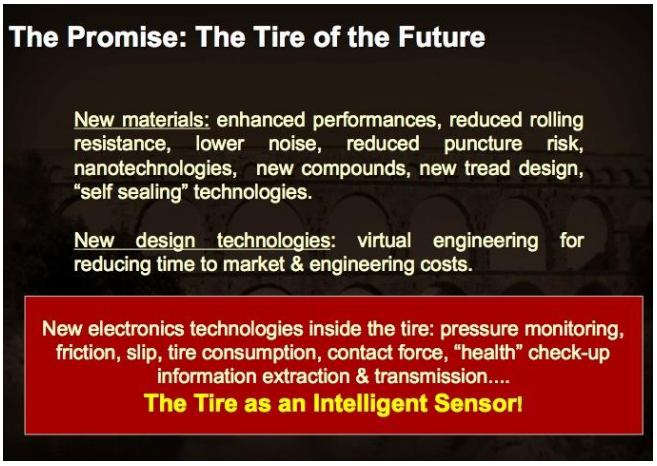
The Promise: The Tire of the Future
New materials: enhanced performances, reduced rolling resistance, lower noise, reduced puncture risk, nanotechnologies, new compounds, new tread design, "self sealing" technologies.
New design technologies: virtual engineering for reducing time to market & engineering costs.
New electronics technologies inside the tire: pressure monitoring, friction, slip, tire consumption, contact force, "health" check-up information extraction & transmission....
The Tire as an Intelligent Sensor!

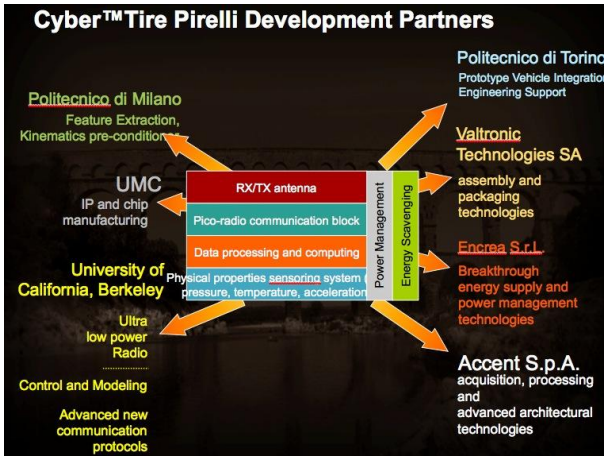
Cyber™Tire Pirelli Development Partners
Politecnico di Torino
Prototype Vehicle Integration
Engineering Support
Politecnico di Milano
Feature Extraction,
Kinematics pre-condition
Valtronic Technologies SA
assembly and packaging technologies
UMC
IP and chip manufacturing
RX/TX antenna
Pico-radio communication block
Data processing and computing
Physical properties sensoring system
pressure, temperature, acceleration
Power Management
Energy Scavenging
Encrea S.r.l.
Breakthrough energy supply and power management technologies
University of California, Berkeley
Ultra low power Radio
Control and Modeling
Advanced new communication protocols
Accent S.p.A.
acquisition, processing and advanced architectural technologies

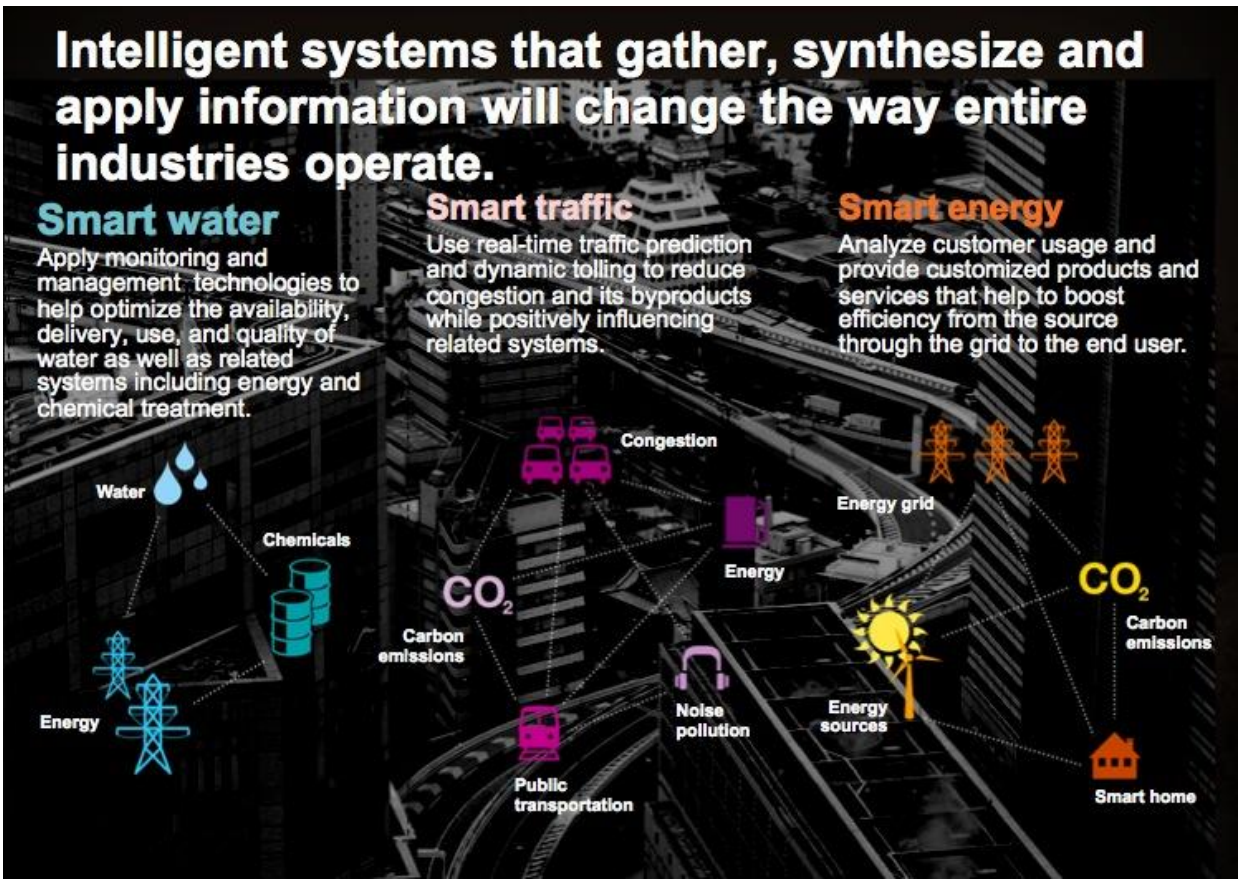
Intelligent systems that gather, synthesize and apply information will change the way entire industries operate.
Smart water
Apply monitoring and management technologies to help optimize the availability, delivery, use, and quality of water as well as related systems including energy and chemical treatment.
Smart traffic
Use real-time traffic prediction and dynamic tolling to reduce congestion and its byproducts while positively influencing related systems.
Smart energy
Analyze customer usage and provide customized products and services that help to boost efficiency from the source through the grid to the end user.
Water
Chemicals
Energy
CO2
Carbon emissions
Congestion
Energy
Noise pollution
Public transportation
Energy grid
Energy sources
CO2
Carbon emissions
Smart home

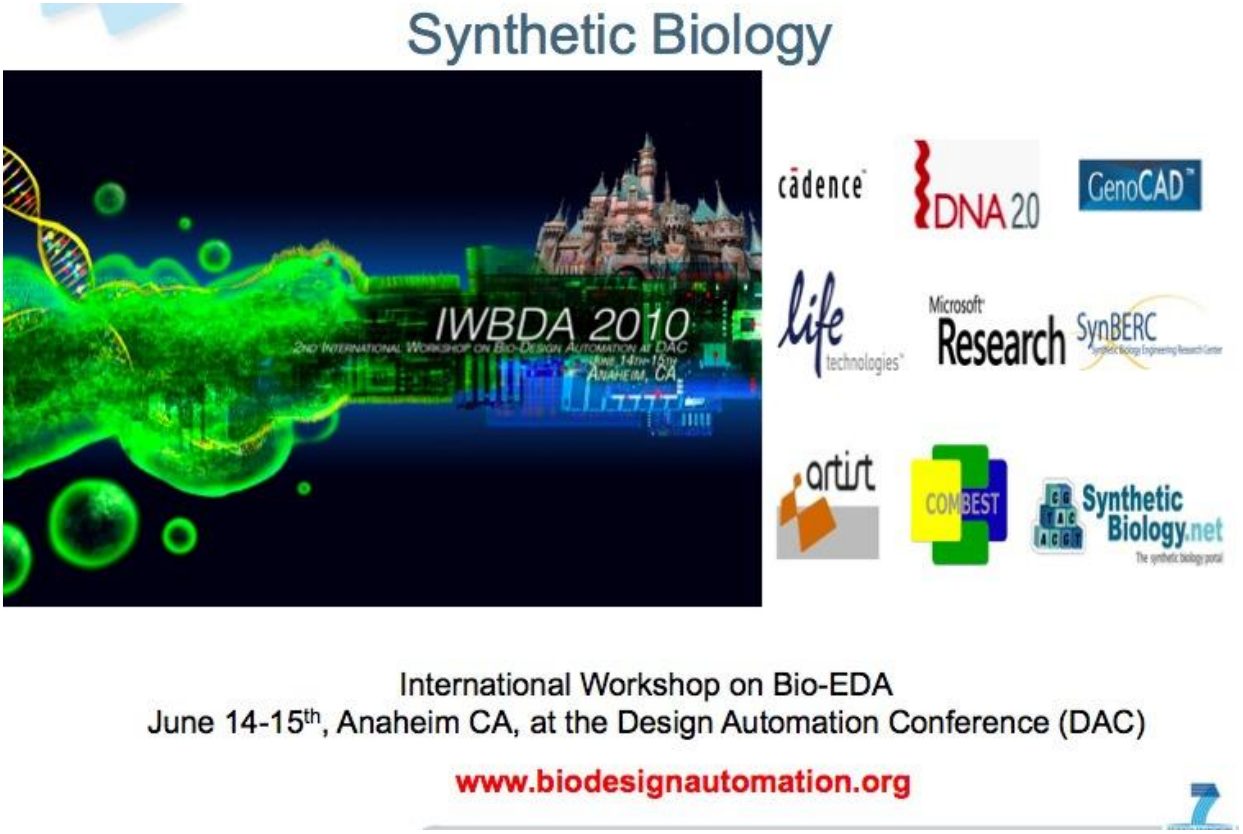

- Synthetic biology seeks, through understanding, to design biological systems and their components to address a host of problems that cannot be solved using naturally-occurring entities
- Enormous potential benefits to medicine, environmental remediation and renewable energy

Applications of Synthetic Biology:

- Energy Crop
  - Water saving
  - No fertilizer
  - Doubled photosynthetic efficiency
- Biodiesel and bio-jet fuel
  - No compromise
  - Fully compatible with existing infrastructure
- Natural product drugs
  - Capture all of the chemistry in nature
  - Construct a microbe that can produce any natural product

1.3.8 Start-ups founded by Artist Design Partners

- AbsInt (http://www.absint.com/), R. Wilhelm, USAAR, area: performance analysis
- SymtaVision, (www.symtavision.com), R. Ernst TUB spin-off, area: real-time system analysis and optimization for automotive and aerospace applications
- UP4ALL (www.uppaal.com) P. Pettersson, W. Yi, K. Larsen, A. David, U. Aalborg, area: formal methods
- Informatik Centrum Dortmund (ICD) (http://www.icd.de/index_eng.html) , P. Marwedel, CEO, area: software for IT systems
- HiQE and Reniance (http://www.hiqe-capital.com/) (Italy) and GreenBox (US), ASV, area: energy efficiency HealthMicro (http://sutisoft.com/hm1_3f/team.html)(US), ASV, area: health devices
- Evidence, G. Buttazzo, RTOS for multiprocessor Systems
- iNoCs, (http://www.inocs.com/), L. Benini, Uni. Bologna, area: Networks on chip.
- Rapita, (http://www.rapitasystems.com/) A. Burns, York, Chair of Board area: tools and services for worst-case execution time analysis (and associated forms of analysis)

1.4 Main events and publications organised

The ArtistDesign High-level Events were intended to gather together the very best world-leading experts from academia and industry, to discuss progress on the state of the art, relevant work directions.

The complete lists of events organised are available here:

- **Schools and seminars**
  http://www.artist-embedded.org/artist/-Schools-and-Seminars,59-.html
  The NoE organized 22 specific schools and seminars.
  This does not include those organised under the earlier Artist2 NoE.
- **International Collaboration events**
  http://www.artist-embedded.org/artist/-International-Collaboration,82-.html
  These are events having a significant impact beyond Europe's borders.
- **Workshops**
  http://www.artist-embedded.org/artist/-Workshops,29-.html
  (some overlap with the international collaboration events)
- **Education**
  http://www.artist-embedded.org/artist/-Transversal-topics-in-Embedded-.html

The most important events were the ARTIST Summer Schools in Europe, China, and South America, as described below.

1.4.1 ARTIST Summer Schools in Europe (2009 2010 2011)

http://www.artist-embedded.org/artist/-ARTIST-Summer-School-in-Europe-.html

http://www.artist-embedded.org/artist/-ARTIST-Summer-School-Europe-2010-.html

http://www.artist-embedded.org/artist/-ARTIST-Summer-School-Europe-2011-.html

The ArtistDesign European Network of Excellence on Embedded Systems Design has organized the 5th through 7th edition of it's highly successful "ARTIST Summer School in Europe" series - funded by the European Commission. These yearly schools on embedded systems design, initiatied in the first ARTIST project, were meant to be exceptional in terms of both breadth of coverage and invited speakers.

These were the largest such schools in the world for embedded systems design (120 participants in 2010), and each year gathered some of the top researchers from Europe, the USA, and Asia.

The schools brought together some of the best lecturers from Europe, USA and China in a 6-day programme, and was a fantastic opportunity for interaction. The 2011 edition was held in beautiful Aix-les-Bains, near Grenoble - France.

1.4.2 ARTIST Summer Schools in China (2009 2010 2011)

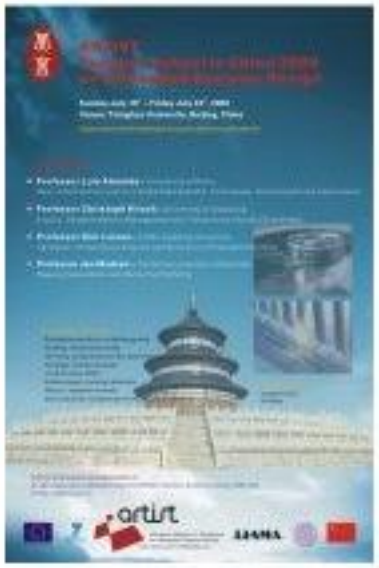

http://www.artist-embedded.org/artist/-ARTIST-Summer-School-in-China-2009-.html

http://www.artist-embedded.org/artist/-ARTIST-Summer-School-in-China-2010-.html

http://www.artist-embedded.org/artist/-ARTIST-Summer-School-in-China-2011-.html

The ArtistDesign European Network of Excellence on Embedded Systems Design has organized the highly successful "ARTIST Summer School in China" series - funded by the European Commission. 2011 was the sixth edition of these yearly schools on embedded systems design.

The school was open in priority to Chinese students. We believe that this will open opportunities for collaboration with Chinese research teams. The school offers a full week consisting of in-depth tutorials on state-of-the-art techniques for the design and analysis of embedded systems given by leading experts.

We aimed to provide a forum for young professors, lecturers, researchers, postgraduates (advanced master and PhD students) working in embedded systems as well as engineers from industry with practical background with the development of embedded systems.

Participants were selected according to their CVs submitted to the organization committee. Programme and Lecturers Top European lecturers from the ArtistDesign European Network of Excellence will provide a world-class programme.

1.4.3 ARTIST Summer Schools in South America (2009 2010 2011)


http://www.artist-embedded.org/artist/-ARTIST-SummerSchool-SouthAmerica-.html (2009)

http://www.artist-embedded.org/artist/-ARTIST-Summer-School-South-America-.html (2010)

The ARTIST Summer School in South America has strengthened the cooperation between Europe and South America in the area of embedded systems, both at educational and research levels. For this purpose, the goal of the school is to provide state-of-the-art courses on embedded systems oriented towards advanced students and young researchers. It should also provide a pleasant atmosphere for research-related discussions among the participants.

1.4.4 Publications

The ArtistDesign community is extremely active in publishing in scientific journals and conferences, and these seem seem to be a reliable measure of integration and building excellence between the partners. In the annual deliverables on Spreading Excellence, there are approximately 80 pages of references to joint papers published in leading journals by the partners over the course of the NoE.

1.5 Impacts on Key Conferences

The ArtistDesign NoE has had a very significant impact on the main international conferences in the area, most specifically :

- ES Week (http://www.esweek.org/) – The vast majority of General Chairs for ES Week in the last 10 years have been members of the ArtistDesign NoE.
- CPS Week (http://www.cpsweek.org/) – CPS Week is the main research conference for Embedded Systems in the USA. The level of interaction between ARTIST members and CPS Week as been fantastic.
- RTSS (http://www.rtss.org/) - RTSS is a major conference in the area, which has been continuously organized by ARTIST members since the start of the NoE.

A fairly complete reference to all active conferences in the area is available on the EMSIG webiste: http://emsig.embedded-systems-portal.org/links/es-conferences/

1.6 Other Impacts

Dissemination efforts were structured arounf the Joint Programme of Activities for Spreading Excellence (JPASE) part of the NoE's workplan. Our actions for Spreading Excellence were at 2 levels:

- *Targeted towards affiliated partners*
  Affiliated partners are not core members in the consortium, but receive support for travelling to ArtistDesign meetings, and actively contribute to the implementation of the Joint Programme of Activities (JPA). These affiliated partners include industrial, SME, academic, and international collaboration affiliates.
- *Targeted towards the scientific and technical community in the large*
  This is achieved mainly bottom-up through the organisation of scientific events, publications, distribution of tools and components, industrial partnerships (not funded by ArtistDesign), education; and through the ArtistDesign web pages.
- *Targeted towards students*
  A particular focus has been placed on the ARTIST Summer Schools this year – with a truly outstanding programme of lecturers, and the innovation of providing the lectures in video form on the Artist website.

Regarding Scientific events, we distinguish between conferences and workshops, schools, and high-level events mainly for International Collaboration.

1.6.1 Tools and Components

The ArtistDesign community plays a leading role in the distribution of software tools and components, on verification/validation tools. Some tools are distributed free of charge, such as UPAAL, IF. Others are commercialised, such as AbsInt, SymTA/S. For many other tools used in the platforms, and shared between the Artist partners, a common dissemination policy has not yet been defined.

1.6.2 Industrial Liaison

ArtistDesign has a wide array of affiliated industrial and SME partners (see the Thematic and Transversal Activity deliverables). Most of these partners have participated in some way in the ArtistDesign technical meetings and the overall effort. There is strong, high-level industry participation through the various Spreading Excellence events organised by ArtistDesign. Our active involvement in the European Technology Platform ARTEMIS also could have a significant and long-term impact.

We believe that the strong involvement of four main ArtistDesign partners in the SPEEDS Integrated Project has a very positive impact on progress in the state of the art, in component-based embedded systems engineering.

1.6.3 Affiliated partners

Affiliated partners are not core members in the consortium, but receive support for travelling to ArtistDesign meetings, and actively contribute to the implementation of the Jointly-executed Programme of Activities (JPA). These affiliated partners include industrial, SME, academic, and international affiliates.

1.6.4 Interaction with the Scientific and Technical Community in the Large

Interaction with these other scientific communities is achieved mainly bottom-up through the organisation of scientific events, publications, distribution of tools and components, industrial partnerships (not funded by ArtistDesign), education; and through the ArtistDesign web pages.

Our sponsoring policy aims specifically at enforcing integration of existing scientific events in the area. This is sought in particular through the Embedded Systems Week (http://www.esweek.org/), in which we play a crucial role.

Another concrete example is our action within the DATE conference (http://www.date-conference.com/), in which we are working to shift the emphasis towards becoming the central European conference on embedded systems design, in collaboration with the ARTEMIS European Technology Platform.

Regarding Scientific events, we distinguish between conferences and workshops, schools, and high-level events mainly for International Collaboration.

The ARTIST community now clearly leads the initiatives for organizing the most significant conferences in the area. In Europe, it has a very strong presence in the DATE conference, which is becoming the main conference on embedded systems within Europe. Over the past 10 years, 9 general chairs of DATE have been leading ARTIST members.

Artist partners are also active members of the ACM's SIGBED, and the IEEE's upcoming Special Interest Group on Embedded Systems currently being set up. Artist members actively work for structuring international events on embedded systems.

1.6.5 International Collaboration

International Collaboration has been one of the central activities pursued within ARTIST since 2003, and is described in detail in this document.

All of the recurring ARTIST International Collaboration events have continued and been expanded within ArtistDesign in 2011. Further details about the schools are available in the section "Organisation of Schools".

1.6.6 Industrial Liaison

ArtistDesign has a wide array of affiliated industrial and SME partners, as described in the deliverables' "Affiliated Partners" sections. Most of these partners participate in some way in the ArtistDesign technical meetings and the overall effort. There is strong, high-level industry participation through the various Spreading Excellence events organised by ArtistDesign.

Our active involvement in the European Technology Platform ARTEMIS also could have a significant and long-term impact. Several ArtistDesign partners, including VERIMAG, BOLOGNA, OFFIS and TU Vienna, are actively involved in the ARTEMIS ETP. The ArtistDesign Strategic Management Board was actively consulted for finalizing the 2011 release of the ARTEMIS Strategic Research Agenda.

In addition, each ArtistDesign partner has an outstanding track record for interaction with industry. Globally, the ArtistDesign consortium has a very strong impact on European R&D in embedded systems. This impact is visible via the achievements in Integrated Projects and STREPs (see below).

1.6.7 Sustainable integration by transfer of selected ARTIST activities to the EMSIG Special Interest Group

Some of the more important initiatives within the Spreading Excellence effort will continue through the newly-created EMSIG (Embedded systems Special Interest Group: http://emsig.embedded-systems-portal.org/) chaired by Peter Marwedel, TU Dortmund. EMSIG is created within EDAA (http://www.edaa.com/) the permanent structure that handles the DATE conference. Most notably, it aims to:

- leverage on the ArtistDesign NoE results
- help to structure the Embedded Systems Design community
- identify areas of future research and development
- interact on a high level with industry, funding agencies, and other world-class structures (eg ARTEMIS)
- support the advancement of state of the art embedded systems design approaches
- facilitate the exchange of ideas and knowledge in embedded systems design
- provide visibility and dissemination (eg: web hosting for events), projects
- support the creation of project consortia
- spread knowledge in embedded systems design
- support the application of embedded systems technology

Shift to the EMSIG website is ensured, by providing links within the NoE website to the equivalent sections on the EMSIG website, and a permanent link on all the NoE website's pages (indicated in red) points to the EMSIG website. We will also make regular announcements to the worldwide Embedded Systems community about EMSIG's activities, and referring to ARTIST NoE as needed.

The ARTIST Summer School is also continuing beyond ARTIST, with support from Nano-Tera (EPFL): http://emsig.embedded-systems-portal.org/links/summer-schools#1027

The integration achieved between tools developed by the ARTIST partners will also provide a basis for continuing integration in the longer term.

1.6.8 Main events organised / publications / impacts on conferences

These include :

- The DATE conference (http://www.date-conference.com/), which has drawn a very large proportion of its General Chairs from the ARTIST NoE's prominent members.
- The CPSWeek conference (http://www.cpsweek.org/), which has always had very tight links with the ARTIST NoE
- Prof. Peter Marwedel's authoritative book "Embedded Systems Design" (http://www.springer.com/engineering/circuits+%26+systems/book/978-94-007-0256-1)
  This is part of a larger book series at Springer on Embedded Systems: http://www.springer.com/series/8563?detailsPage=titles

1.6.9 Links with ARTEMISIA

ArtistDesign has strong links to ARTEMIS, through:

- Representation on the **ARTEMIS Industry Association Steering Board**:
  - Joseph Sifakis was the initial CNRS representative
  - Luca Beninni is the University of Bologna representative
- Partner membership in **ARTEMIS "B"** (Research Organisations & Universities) http://www.artemisia-association.org/member_status
  - Arne Skou is the Aalborg University representative
  - Denis Platter is the CEA representative
  - Joseph Sifakis is the CNRS-Verimag representative
  - Boudewijn Haverkort is the Embedded Systems Institute representative
  - Rudy Lauwereins is the IMEC representative
  - Jean-Pierre Banâtre is the INRIA representative
  - Eduardo Tovar is the Instituto Superior de Engenharia do Porto representative (Instituto Politécnico do Porto in ArtistDesign)
  - Gunnar Landgren is the KTH representative
  - Bernhard Josko is the OFFIS representative
  - Jan Madsen is the TU Denmark representative
  - José Carlos Gómez Sal is the University of Cantabria representative
  - Luca Benini is the University of Bologna representative
  - Farid Ouabdesselam is the Université Joseph Fourier representative

- Strong *informal* links. For example, the ArtistDesign Strategic Management Board was asked to review and comment on the latest edition of the Strategic Research Agenda, published in 2011.
- Strong representation by ArtistDesign partners in ARTEMIS projects,

1.6.10 ArtistDesign Web Portal

Objectives and Background Information

The ArtistDesign Web Portal was a major tool for Spreading Excellence within the Embedded Systems Community. Its aim was rather ambitious: to be the focal point of reference for events and announcements of interest to the embedded systems community.

The web portal disseminates information about contacts (ArtistDesign core and affiliated partners), the ArtistDesign JPA activities, as well a fairly thorough set of links to sites of interest to the embedded systems community.

As can be seen, a great deal of effort has been put into the web site, both for ergonomics / graphical quality, as for the contents.

The web site includes several features that help keep it coherent and up to date:

Authorised users (principally, the ArtistDesign partners) can access the back end of the site to modify and update information directly. The changes are immediately visible on the site, which greatly streamlines the updating process.

It's possible to track changes and go back to previous versions of individual web pages.

Events are automatically sorted by date, and transferred to 'Past Events'. When appropriate.

Structural information (hierarchy of pages) was maintained automatically.

Ergnomics are set for the entire site. The "look and feel" of the site was always homogeneous thoughout the site. It's possible to change these ergonomics, and these changes are applied homogeneously throughout the site, via automated machanisms.

Structure

The structure of the ArtistDesign web site was visible on the Site Map: http://www.artist-embedded.org/artist/spip.php?page=plan).