IST-004527 ARTIST2
Network of Excellence
on Embedded Systems Design

Publishable Final Activity Report

Bruno Bouyssounouse (Editor)
UJF/Verimag Laboratory
Artist2 Technical Coordinator

Artist2 NoE Consortium

Policy Objective (abstract)

The ARTIST2 NoE is structured around a set of intra and inter-cluster research activities on cutting-edge topics in embedded systems design, reflecting the following decomposition of the embedded systems design flow. This forms the essential ingredient of the NoE, which has strongly motivated the participating research teams.

Artist2 NoE Participants

For a complete description including web links, see:
http://www.artist-embedded.org/artist/-Core-Partners-.html

Scientific Coordinator: Joseph Sifakis
Tel: +33 4 56 52 03 51
Joseph.Sifakis@imag.fr

Technical Coordinator: Bruno Bouyssounouse
Tel: +33 4 56 52 03 68
Bruno.Bouyssounouse@imag.fr

Mailing address: Verimag Laboratory - Centre Equation - 2, ave de Vignate - 38610 Gières - France

Partic N°	Participant name	Person in charge
1	Caisse des Dépots et Consignations (France)	Jean-François Forté
2	UJF / Verimag (France)	Joseph Sifakis
3	RWTH Aachen (Germany)	Rainer Leupers
4	BRICS – Aalborg University (Denmark)	Kim Larsen
5	AbsInt GmbH (Germany)	Christian Ferdinand
6	University of Aveiro (Portugal)	Luis Almeida
7	Universidad de Cantabria (Spain)	Michael Gonzalez
8	CEA –LIST (France)	Francois Terrier
9	CFV, Université de Liège (Belgium)	Pierre Wolper
10	Czech Technical University (Czech Rep.)	Zdenek Hanzalek
11	Dortmund University (Germany)	Peter Marwedel
12	TU Denmark (Denmark)	Jan Madsen
13	ETH Zurich (Switzerland)	Lothar Thiele
14	France Telecom R&D (France)	Jacques Pulou
15	INRIA (France)	Alain Girault
16	KTH (Sweden)	Martin Törngren
17	Linköping U. (Sweden)	Petru Eles
18	CNRS / Laboratoire LSV (France)	Philippe Schnoebelen
19	Lund University (Sweden)	Karl Erik Arzen
20	U. Mälardalen (Sweden)	Björn Lisper

Partic N°	Participant name	Person in charge
21	Kuratorium OFFIS e. V. (Germany)	Werner Damm
22	PARADES (Italy)	Alberto Sangiovanni
24	UP Madrid (Spain)	Alejandro Alonso
25	Saarland University (Germany)	Reinhard Wilhelm
27	TU of Eindhoven (Netherlands)	Jeroen Voeten
28	TU Vienna (Austria)	Hermann Kopetz
29	TU Braunschweig (Germany)	Rolf Ernst
30	University of Twente (Netherlands)	Ed Brinksma
31	University of Bologna (Italy)	Luca Benini
32	Uppsala University (Sweden)	Bengt Jonsson
33	Universidad Polytechnical de Valencia (Spain)	Alfons Crespo
34	University of York (UK)	Alan Burns
35	Polytechnic Institute of Porto (Portugal)	Eduardo Tovar
36	EPFL (Switzerland)	Tom Henzinger
37	Scuola Sant'Anna – Pisa (Italy)	Giorgio Buttazzo
38	Ace (Netherland)	Joseph van Vlijmen
39	Tidorum (Finland)	Niklas Holsti
40	TU Kaiserslautern (Germany)	Gerhard Fohler
41	TU Berlin (Germany)	Sabine Glesner

The following partners have withdrawn from the NoE:
Partner 23 – University of Pavia (Italy) Partner 26 – ST Microelectronics (France)

Table of Contents

1. High-Level Objectives

1.1 Strengthening Scientific and Technological Excellence for Embedded Systems Design

ARTIST2 has implemented an international and interdisciplinary fusion of effort to foster the emergence of the European research community on embedded systems design. This ongoing interdisciplinary effort in research is mandatory to fully establish embedded systems design as a discipline combining competencies from electrical engineering, computer science, applied mathematics, and control theory. The ambition is to compete on the same level as equivalent centres in the USA (Berkeley, Stanford, MIT, Carnegie Mellon), for both the production and transfer of knowledge and competencies, and for the impact on industrial innovation.

For example, in Berkeley, the Centre for Hybrid and Embedded Software Systems (CHESS), funded by a large ITR grant of NSF, includes more than ten faculty members in all areas of the Department of Electrical Engineering and Computer Sciences from hybrid control systems, to formal methods, from chip architectures to software systems, from compilers to operating systems. The Centre is connected with other institutions such as Vanderbilt University and has a strong industrial backing from the aeronautics, automotive, computers, components and EDA industry.

At the same institution, the MARCO Gigascale System Research Centre (GSRC) involves seven faculty members from Berkeley and twenty from other institutions to develop new design methodologies with strong emphasis placed on reconfigurable and heterogeneous platforms.

One of the main objectives of the ARTIST2 Network of Excellence has been to gather together the best European teams from the composing disciplines, to forge a scientific community. This objective has been achieved by integration around a Joint Programme of Activities, aiming to create critical mass from selected European teams. Integration occurred at two levels:

- Integration within clusters, corresponding to essential topics in the area of embedded systems design. Before the start of the NoE, efforts on the identified topics are fragmented, and there was no European research team that could gather the sufficient critical mass needed. The integration of a topic is a first step towards integrating the area as a whole.
- Integration between clusters topics to create the multi-disciplinary community which now plays a major role in driving research in the area of embedded systems design. This was achieved through integration activities that will bring together teams from different clusters.

Building the embedded systems design scientific community has been an ambitious programme, and continues through the ArtistDesign Network of Excellence. ARTIST2 has built on the achievements and experience from the ARTIST1 FP5 Accompanying Measure (http://www.artist-embedded.org/) on Advanced Real-Time Systems. Beyond ARTIST1, there was a remarkably strong willingness within the composing scientific communities to establish this new area – by accelerating the ongoing convergence between them. In the USA, this convergence has been strongly supported by funding agencies (e.g.: DARPA and NSF), through the creation of dedicated research centres and R&D programmes, by launching and dedicated conferences (e.g.: EmSoft), special interest groups (e.g.: SIGBED of the ACM http://www.acm.org/sigbed/) and journals (e.g.: Transactions in Embedded Computing Systems [TECS] http://www.acm.org/tecs/).

Even if a number of the core partners in ARTIST2 were initially from different disciplines, many had already interacted, due to their participation in a multitude of R&D projects – sponsored by

the EC under FP5, Esprit and Eureka. Within these projects, many shared common industrial partners.

The partners' international standing, research and teaching programmes in the field, the technologies they have developed and possess, and their leading presence in international scientific events prove their excellence. The Joint Programme of Activities integrated the partner institutions mainly by promoting close collaboration and massive researcher exchanges between partners and thus start an ARTIST culture within the network.

1.2 Spreading the Excellence in Embedded Systems Design

The core partners represented approximately 140 senior researchers. Including the affiliated partners, this reached approximately nearly 450 researchers in all. With this number of excellent people working on the same goals, the visibility of the European research effort in embedded systems design is now worldwide. This has progressively created a European embedded systems design community, and spread 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 ARTIST2 knowledge in education and training. Activities such as joint PhDs between partners, courseware, textbook publications, summer schools, and seminars, have all served to attract students and young researchers to our research field, without gender discrimination.

A specific effort has been devoted to spreading excellence to industry. All the core partners have strong and lasting collaborations with major industrial players in the area. Partners have seen the evolution of industrial know-how and techniques, through the integration of state of the art results. Concretely, this transfer to industry was implemented through shared PhDs with affiliated industrial partners, by opening the NoE's platforms to industry, and through a strong participation of the partners in FP6/FP7 Integrated Projects, STREPS, etc.

1.3 A Lasting ARTIST Network of Excellence

One of the main objectives for the consortium was to successfully implement lasting integration in the area. This was achieved by setting up the appropriate instruments and structure.

ARTIST2 established a common infrastructure (web portal, communication tools, and video) to enhance interaction and collaboration between partners, and also with the industrial and research communities at large. ARTIST2 was able capitalize on knowledge through a set of instruments, to build a common reference point for the development of research and training of new researchers.

- Platforms bringing together tools and competencies, and making them available to research, education and industry.
- Graduate and Summer Schools led by ARTIST2 partners have been created.
- Web Portal, acting as a repository of knowledge in the area, including courseware, information about standards, methods and tools, research publications and results. This web portal will be made available within the NoE core and affiliated partners, and also to other parties according to modalities to be defined.

ARTIST2's four-year programme has had a strong commitment to integration and sustainability. This commitment has guided the definition of the Joint Programme of Activities, and the vision for financial autonomy. To achieve the aims, the support from the EC was a very small proportion of the overall investment by the core partners.

2. Structure of the Research Effort

2.1 Overview

The ARTIST2 NoE is structured around a set of intra and inter-cluster research activities on cutting-edge topics in embedded systems design, reflecting the following decomposition of the embedded systems design flow. This forms the essential ingredient of the NoE, which has strongly motivated the participating research teams.

This strong technical focus has led to very specific technical aims composing the strategy for using and disseminating knowledge created within the NoE.

Our strategy for dissemination and use of knowledge is at 3 levels:

- *Targeted towards the NoE partners*
  Artist2 core partners are the heaviest users of all the dissemination actions described in this document – both as organisers and as participants.
- *Targeted towards affiliated partners*
  Affiliated partners are not core members in the consortium, but receive support for travelling to Artist2 meetings, and actively contribute to the implementation of the Joint Programme of Activities (JPA). These affiliated partners include industrial, SME, academic, and international 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 Artist2), education; and through the Artist2 web pages.

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

This design flow is 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.

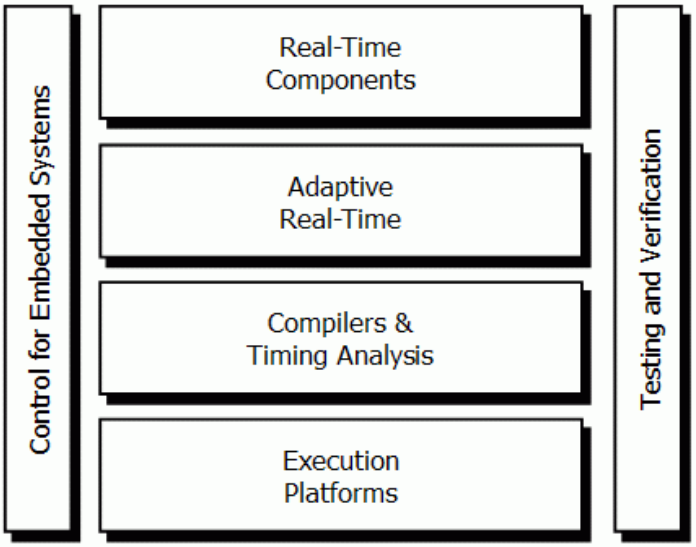

Real-Time Components: The development of a general framework for component-based engineering of complex heterogeneous embedded systems is a grand challenge which spans many research problems. A key characteristic of component-based embedded systems is heterogeneity of component models. This heterogeneity concerns different execution models (synchronous, asynchronous, vs. timed), communication models (synchronous vs. asynchronous), as well as different scheduling paradigms. Technology must be provided to allow designing heterogeneous embedded systems from diverse types of components, and allow predicting and optimizing functional and non-functional properties of the designed systems.

Adaptive Real-time: This is a more recent approach to embedded systems design, where temporal constraints can be relaxed, which allows optimized use of resources. This includes applications – where managing the Quality of Service (QoS) is essential, such as telecommunication systems, multi-media, and wide-area networked applications. In this relatively new area, there is a recognized lack of design theory and tools.

Compilers and Timing Analysis: Once the application software has been developed, using the above, the system must be implemented on a given target platform. Compilation tools and their associated technologies play a fundamental role for automating this process. For the implementation of embedded systems, we need tools capable of combining platform independent software and a description of the target platform, to generate an executable code having the desired properties related to use of such resources as memory, power, energy, network bandwidth, and computation time). Resource-aware compilation requires the use of Timing Analysis tools to estimate the execution times of embedded software on a given platform.

Execution Platforms: This topic is strongly linked to the compilation and implementation of embedded systems. For a given application, it is important to have the technology, methods and tools to make rational choices about the platform and the design used, before proceeding to final implementation. Research in Execution Platforms targets the development of the theoretical and practical tools for modelling the dynamic behaviour of application software for a given platform. This is a new area of research, which will allow greater flexibility in designing optimal embedded systems.

Testing and Verification: This is transversal topic, which interacts with all the other topics in embedded systems design. It aims to ensure that the different design steps meet given properties, as well as the overall correctness of the implementation. This is a very active research topic, with results at different levels of the design process. The current challenge is in achieving on overall approach for testing and verification, focussing on two important aspects.

First is the Verification and Testing of real-time properties, to ensure that hard real-time constraints or quality of service constraints are met. Second is for Verification of Security Properties, where identification of gaps in security is desired.

Control for Embedded Systems: Embedded systems are deployed in the real world, and are often reactive to it. This interaction with the environment is intrinsic to the service provided. A large proportion of embedded systems can be considered to be controllers. On the other hand, most automated control applications will be implemented as embedded components. Thus, it is essential that work on joining control theory and embedded systems be included in the ARTIST2 NoE.

2.2 Detailed View of the Research Effort

2.2.1 Real-Time Components Cluster

This cluster is composed of the following activities:

Component Modelling and Verification (Platform)
Led by Susanne Graf (UJF/ VERIMAG)

Partner teams (leaders): Susanne Graf (VERIMAG), Saddek Bensalem (Verimag), Michael Perin (Verimag), Laurent Mounier (Verimag), Olivier Constant (Verimag), Sébastien Gérard (CEA LIST), Francois Terrier (CEA LIST), Jacques Pulou (France Telecom R&D), Thierry Coupaye (France Telecom R&D), Noël Plouzeau (INRIA), Bernhard Josko (OFFIS), Alberto Sangiovanni-Vincentelli (PARADES), Wang Yi (Uppsala), Bengt Jonsson (Uppsala University).

Affiliated teams (leaders): Julio Medina (U. of Cantabria), David Lesens (EADS), Alain Leguennec (Esterel Technologies), Veronique Fabre (Thales), Martin Torngren (KTH).

Overview: Before ARTIST started, UML was becoming a standard for model-based development, also in the context of real-time and embedded systems, even if it was lacking a number of concepts needed for this purpose and supporting validation tools. In the context of real-time embedded systems, there existed a number of UML based CASE tools (e.g., Artisan, Rhapsody, RoseRT, and TAU) and a large number of analysis and validation tools, mostly coming from academia. With a few exceptions, they were dedicated to specific profiles taking into account a small subset of UML and are weakly integrated in the development flow.

Several of the platform participants had already started considerable efforts for integrating analysis and validation into the development flow --- in particular in the framework of IST projects AIT-WOODES (CEA: Accord Methodology and tool support, OFFIS: verification tool for UML in Rhapsody), OMEGA (VERIMAG: IF verification tool for real-time UML, OFFIS: verification tool for UML), and Metropolis (PARADES: UML platform).

Work in Year 1

The main objective of the first year of the project was to obtain an inventory of potentially interesting tools, possibly to do some initial developments within these tools towards a possible integration and finally to define a concrete vision of the ARTIST platform for component-based design and validation.

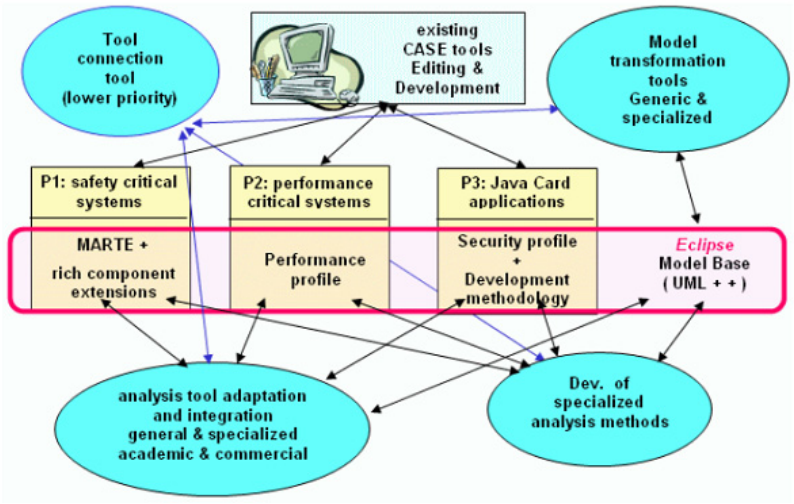

Planned platform architecture

We had chosen the option to first connect a restricted set of model-based analysis and validation tools with the help of tools implementing UML compatible model transformation technology and possibly – if this turns out to be useful – tools allowing to generate complex functionalities from basic ones by means of abstract specifications. The set of participating tools is always to be considered preliminary; new tools were expected to join the platform over time.

During the first year of ARTIST, we have done only a limited amount of integration. The main progress was on individual components for these platforms, whose description can be found in the year 1 deliverables.

Work in Year 2

The outcomes and achievements of the second year have been structured into five main topics, enumerated below.

- *Semantic foundations for modelling languages and frameworks*
  An important issue for our platform is achieving tool chains for related profiles by mapping them to a small set of semantic level formalisms used in validation and code generation tool chains.
- *Platform for the analysis of safety critical embedded systems*
  These concern back-end tool chains, starting from one of the envisaged semantic level formats and integrating validation and code generation tools.
- *Platform for the analysis of performance critical systems*
  The aim is the integration of performance evaluation and formal verification in requirement and design activities.
- *Platform for the certification of smart-card applications*
  This has not progressed according to the plans that included the definition of a UML profile for security properties, Rather than working on the profile during the first year, it was decided to focus on the validation engine.
- *Generic validation technology for non functional properties and component systems*
  The development of new verification techniques is not the primary goal of the component platform. The focus here is on the connection of existing verification tools to the modelling languages considered in the platform.

Work in Year 3

-- Modelling languages and semantic frameworks and their implementations --

- The MARTE UML profile for modelling real-time systems has been finalised and partially implemented. Also work on a complementary profile for fault tolerance and safety requirements has been continued, as well as the work on an executable UML profile.
- In the SPEEDS project, a *rich component model* (called HRC, standing for Heterogeneous Rich Components) has been defined and implemented as a standalone metamodel.
- The BIP framework developed by VERIMAG has been enriched with hierarchical connectors and a notion of component encapsulation; BIP connectors are now part of the HRC metamodel.
- Metropolis II that is centered on the coordination of components has been finalised. The Metropolis meta-model concepts have been provided as input to the HRC modelling effort in SPEEDS.

-- Platform for the analysis of safety critical embedded systems --

The essential effort was on dedicated analysis tools implemented in tool chains, integrating validation and code generation tools. The main work on the mappings from user level profiles to semantic level formalisms has started towards the end of year 3 for the MARTE profile and and a bit later for HRC.

-- Platform for the analysis of performance-critical systems --

This platform consisted of two unrelated parts: A tool chain allowing the integration of a state-of-the-art performance simulator into the design flow of service oriented systems and a simulation-based approach for evaluating energy related properties in sensor networks.

-- Platform for the certification of smart-card applications --

We have defined a methodology for the certification of smart-applications according to the International standard known as the *Comon Criteria* (CC) for security and developed tools for its support. The methodology uses formal methods to reach the highest level of certification (Evaluation Assurance Level 7+): full formal development. All the tools are integrated in the Eclipse environment and support the entire development from UML specifications to the verdict of the verification tools and certification documents. The TLFIT tool helps to produce the documentation required for certification, it eases the traceability of the security requirements from their expression in natural language to the specification of the formal security policy and its final implementation.

-- Transversal validation technology for platforms --

Work has focused on a set of validation methods and tools which are developed in the context of one of the platforms.

- A prototype tool *CATS*, www.timestool.com/cats, for compositional timing and performance analysis has been developed for systems modeled using timed automata and the real time calculus developed at EPFL. The tool combines timing and performance analysis [HP07].
- The symbolic execution kernel *Agatha* allows exploiting UML models in order to generate requirement test cases.
- We have designed and partly implemented a tool which verifies absence of deadlocks for *BIP* specifications (http://www-verimag.imag.fr/~async/BIP/bip.html). We have developed abstract criteria based on dependency graphs. Finally, we have considered methods exploiting the global Petrinet defined by a BIP system for checking deadlock freedom for BIP component systems.
- *Uppaal* has been extended by methods for directed model-ckecking [KD*07, KD*07b], techniques for validation of hybrid systems with large discrete spaces have been obtained by combining techniques [DD*07], [Seg07]. Cyclic timed automata have been used as a bridge between timed automat and timed event streams allowing the combined use of analysis techniques based on them.

Final Results

We collaborated on two transversal topics and on three platforms.

The transversal topics are related model representations shared by tools and to backend engines:

- Topic 1: Modelling languages and semantic frameworks and their implementations
- Topic 5: Tranversal validation technology (partly shared with T&V cluster)

Each of the platforms provides one or several loosely coupled tool chains. The common denominator of each tool chains is the application domain, that is, the kind of problems to be tackled and also the modelling concepts used. Indeed, the tools of the platforms generally either share a common modelling framework, or are based on complementary (respectively similar) frameworks which show potential for convergence or integration (see also Figure 1, in Section 2.1):

- Topic 2: A platform for the development of safety-critical embedded systems
- Topic 3: A platform for the analysis of performance critical service-based systems
- Topic 4: A platform for the certification of smart-card applications

The first platform is mainly about generic techniques whereas the two other platforms target more specific application domains. A long term goal is reaching a state where they are indeed instances of a common platform. However, since this approach needs a significant amount of dedicated resources, it is outside the scope of Artist.

Development of UML for Real-time Embedded Systems (Cluster Integration)
Led by Sébastien Gérard (CEA - LIST)

Partner teams (leaders):: Sébastien Gérard – CEA, Susanne Graf – VERIMAG, Julio Medina - Cantabria university

Affiliated teams (leaders): Ivica Crnkovic – MdH, Stefan van Baelen - K.U. Leuven, Bernhard Josko – OFFIS, Gilbert Edelin and Dr. Laurent Rioux – Thalès Research and Technology, Matthias Grochtmann – DaimlerChrysler, Henrik Lönn – Volvo

Overview: Since the adoption of the UML standard and its new advanced release UML2, this modelling language has been used for development of a large number of time-critical and resource-critical systems. Based on this experience, a consensus has emerged that, while a useful tool, UML is lacking in some key areas that are of particular concern to real-time and embedded system designers and developers. In particular, it was noticed that first the lack of quantifiable notions of time and resources was an impediment to its broader use in the real-time and embedded domain. Second, the need for rigorous semantics definition is also a mandatory requirement for a widespread usage of the UML for RT/E systems development. And third, specific constructs were required to build models using artefacts related the real-time operating system level such as task and semaphore.

Fortunately, and contrary to an often expressed opinion, it was discovered that UML had all the requisite mechanisms for addressing these issues, in particular through its extensibility facilities. This made the job much easier, since it was unnecessary to add new fundamental modelling concepts to UML – so-called “heavyweight” extensions. Consequently, the job consisted in defining a standard way of using these capabilities to represent concepts and practices from the real-time and embedded domain.

Hence, this specification of a UML™ profile adds capabilities on the one hand for modelling Real Time and Embedded Systems (RT/ES), and on the other hand for analyzing schedulability and performance properties of UML specifications. This new profile is intended to replace the existing UML Profile for Schedulability, Performance and Time [UML profile for Schedulability, Performance, and Time, version 1.1., formal/05-01-02, 2005]. This extension, called the MARTE profile, should address specification, design, and verification stages of the development cycle of RT/ES. It wants to address the two branches of the V cycle, i.e. design and validation& verification. Modelling capabilities have to ensure both hardware and software aspects of RT/ES in order to improve communication/exchange between developers. It has

also to foster the construction of models that may be used to make quantitative analysis regarding hardware and software characteristics. Finally, it should enable interoperability between developments tools used all along the development process.

Work in Year 1

The first year of this activity has been dedicated to firstly influence on the writing of the request for proposal (RFP) of the new UML profile for real-time and embedded systems. This RFP expresses all the requirements the new standard will have to satisfy. The RFP, document referenced at OMG web server as realtime/05-02-06 (UML Profile for Modelling and Analysis of Real-Time and Embedded systems (MARTE) RFP)) has been voted and accepted in the context of the Real-time, Embedded, and Specialized Systems (RTESS) Platform Task Force in February 2005: UML Profile for Modelling and Analysis of Real-Time and Embedded systems (MARTE) RFP , realtime/05-02-06, http://www.omg.org/cgi-bin/doc?realtime/05-02-06.

Within the second half year period, the job consisted in both following action (main part of this work has been performed within the French CARROLL-Protes project):

– To setup an OMG submitter team in order to answer to the RFP. The team that has been organized is called the ProMARTE team: www.promarte.org. This team consists of the main companies (end users and tool providers) involved in this aspect at the OMG. It is composed of: Artisan, Carlton university, CEA, IBM, I-Logix, INRIA, Looked-Martin, Thales, Tri-Pacific.

– To write the initial submission of the ProMARTE team that has been delivered in November 2005: Joint UML Profile for MARTE Initial Submission, realtime/05-11-01, http://www.omg.org/cgi-bin/doc?realtime/2005-11-01

Within this first year, in the context of the Omega project, Verimag aimed at the definition of an UML profile appropriate for real-time embedded systems based on the existing SPT profile. The extension done in Omega introduces a notion of "observer" and emphasizes the importance of capturing the relevant events which make reference to the system at execution and is used to capture its dynamic properties.

Work in Year 2

-- A consolidated architecture for the MARTE profile --

The MARTE profile architecture model consists of three main packages:

– The Time and Concurrent Resource Modelling package (TCRM); it defines basic model constructs for time and resource, especially concurrent resources. These foundational concepts are then refined in both of the following package in order to fit with both modelling and analyzing concerns.

– The Real-Time and Embedded application Modelling package (RTEAM); it enables modelling of RT/E application. It concerns mainly defining high-level model constructs to depict real-time and embedded features of application, and to enable the description of execution platforms, software as well as hardware.

– The Real-Time and Embedded application Analysis; it provides a generic support for analyzing annotated models. This generic framework is also refined in order to cope with schedulability and performance analysis. It is also expected that the generic framework for analysis will be specialized/extended to support other kind of quantitative analysis, such as power consumption, memory use or reliability.

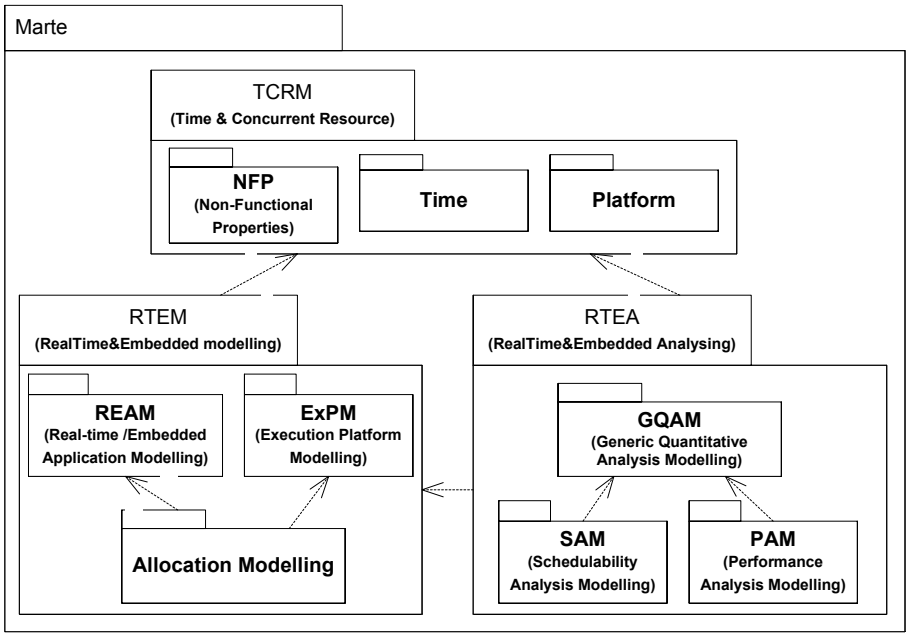

Figure 1. Architecture of the MARTE profile at the end of year 2.

Difficulty: Dissemination issue due to privacy rule of an ongoing work at OMG
Actually, OMG ongoing work performed in OMG consortium such as the ProMarte one is considered to be private until final vote. So, the only available documents related to MARTE in the year 2 were the initial submission that provides only an outline of the proposal; it does not go into the details of the proposed concepts. So, only members of the ProMarte consortium can access the full information of the standard, i.e. CEA (the leader of the activity), INRIA, Thales and Cantabria University.

Work in Year 3

-- Harmonization of the MARTE standard --
(CEA, Cantabria University and Thalès Research and Technology)
Work performed within this activity consisted in making consistent the whole specification which is made of several dependant parts. The main goal was then to harmonize/align all sub-profiles contained within the MARTE specification. As for example, both following specific profiles dedicated to platform modelling, the Hardware Resource Model (HRM) and the Software Resource Model (SRM) have been aligned with the more generic profile contains in the MARTE foundation package, the Generic Resource Platform (GRM) subprofile. This latter defines a set of standard concepts dedicated to model system-level computing platform. All the jobs performed within this activity led to a new MARTE architecture as shown on the following figure.

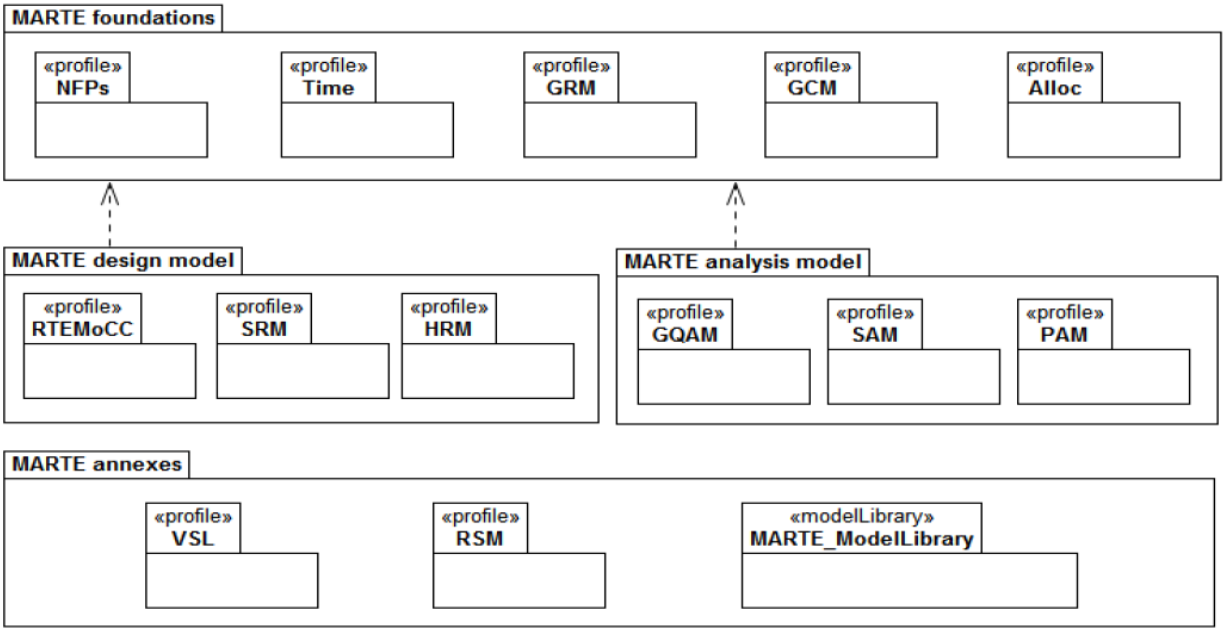

Figure 2. Architecture of the MARTE profile at the end of year 3 (NFPs=Non-Functional Properties, GRM=Generic Resource Modeling, GCM=Generic Component Model, Alloc=Allocation Modeling, RTEMoCC=Real-Time Embedded Model of Communication and Computation, SRM=Software Resource Modeling, HRM=Hardware Resource Modeling, GQAM=Generic Quantitative Analysis Modeling, SAM=Schedulmability Analysis Modeling, PAM=Performance Analysis Modeling, VSL=Value Specification language and RSM=Repetitive Structure Modeling)

Debugging of the MARTE standard (CEA, Cantabria University and Thalès Research and Technology)
Within this period, we have continued to experiment with MARTE in different case studies. All these experiments were used to debug MARTE and hence have contributed a lot to improve its quality and soundness.

Dissemination
The end of this period was also dedicated to build the first materials needed to disseminate MARTE among industry and academics. A web site (hosted by the OMG consortium) has been set up. Among other, this web site gathers all the papers and events related to MARTE, and a very detailed tutorial: http://www.omgmarte.org. Wihtin this web site, you may also have information about current implementation of this specification.

Final Results

-- Experiments of the MARTE standard (CEA and Thalès Research and Technology) --
Within this period, we have continued to experiment the usage of MARTE in two ways. Firstly, we have used MARTE to build exemples of systems in order to assess the usage of MARTE for modelling real-time and embedded features. Within this period, we have implemented the profile within two UML2 tools. CEA did it for the Papyrus open-source tool and Thales did it for the RSA (from IBM) tool. This latter implementation is also available in open-source and maybe downloaded on the OMG's web site for MARTE. In addition, both CEA and Thales have also started to experiment the usage of MARTE for annotating UML models in order to performed schedulability analysis. CEA has then implement first prototype of gateway between papyrus and both schedulability analysis tools, SymTA/S and MAST. And Thales has implemented two other gateways between RSA for the UML modling aspect and RapidRMA and Cheddar for the

schedulability analysis aspect.
www.omgmarte.org and www.papyrusuml.org

-- Raising issues and proposition of resolutions --
In consequence of the previous activity, we have discovered issues within the MARTE specification. Firstly, we have then raised officially all these issues using the OMG issue web service. Secondly, as we were also directly involved in the OMG's finalization task force decicated to MARTE, we have also proposed the adequate resolutions to solve all of these issues. Let's notice that OMG has accepted in June 2008 the finalization task force report we have contributed to. Finally, following figures illustrate the final architecture of the MARTE profile.

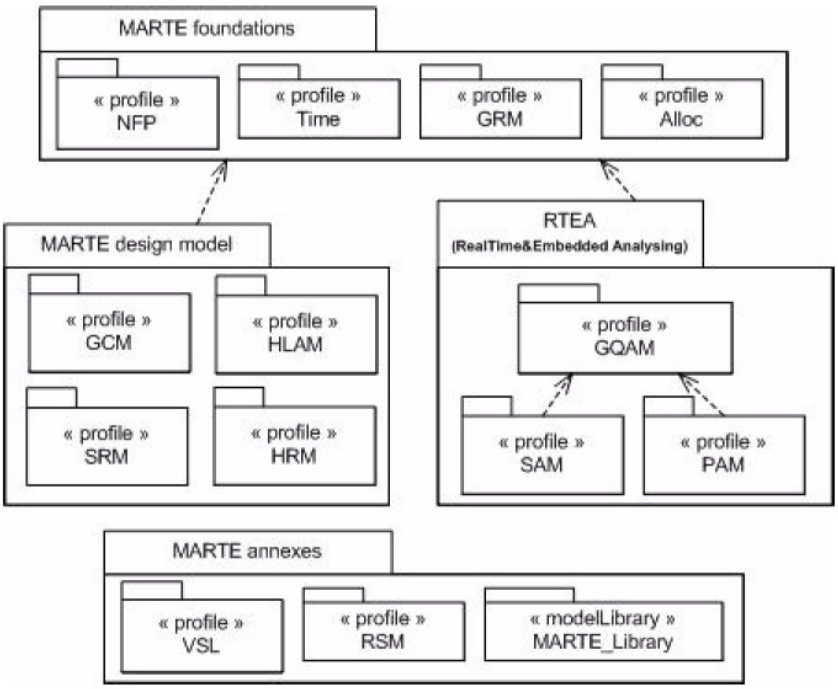

Figure 3. Final architecture of the MARTE profile (NFP = Non-Functional Properties, GRM = Generic Resource Modeling, GCM = Generic Component Model, Alloc = Allocation Modeling, HLAM = High-Level Application Modeling, SRM = Software Resource Modeling, HRM = Hardware Resource Modeling, GQAM = Generic Quantitative Analysis Modeling, SAM = Schedulmability Analysis Modeling, PAM = Performance Analysis Modeling, VSL = Value Specification language and RSM = Repetitive Structure Modeling)

-- Disseminations and relationship with industry --
From its adoption by the OMG in June 2007, in addition to previous acticities that consists firstlsy in elaborating examples to prove the validity of the specification and find possible issues, and secondly in improving the quality of the specification by solving the rasied issues, one important task has been to promote as much as possible this new standard through production of tutorials and tools supporting the specification. This specific effort of dissemination of MARTE is specillay concretized through the new FP 7 project ADAMS. It aims at promoting the usage of the MARTE standard for the development of real-time and embedded systems using both model and component design paradigms. www.adams-project.org.

In the Persiform project, Verimag and FTRD had defined a UML profile for modelling for performance adopting performance concepts from a major performance analysis tool (Hyperformix workbench) [CWG08]. In the OpenEmBeDD project, Verimag and Thales work on

an adaptation of the developed tool chain to a subset of MARTE. A complete adaptation however would need enriching MARTE which should be envisaged for a later point of time.

Component-based Design of Heterogeneous Systems (Cluster Integration)
Led by Bengt Jonsson (Uppsala)

Partner teams (leaders): Bengt Jonsson (Uppsala), Francois Terrier – CEA (France), Tom Henzinger – EPFL (Switzerland), Albert Benveniste – INRIA (France), Jean-Marc Jézéquel - Inria (France), Werner Damm - OFFIS (Germany), Alberto Sangiovanni-Vincentelli - PARADES (Italy), Paul Caspi – Verimag (France), Joseph Sifakis – Verimag (France), Hermann Kopetz - TU Vienna (Austria), Jacques Pulou (FTRD, France)

Affiliated teams (leaders): Anders Ravn – Aalborg (Denmark), Peter Eriksson - ABB Automation Technology (Sweden), Bernhard Steffen - Dortmund University (Germany), Ivica Crnkovic – MdH (Sweden), Frédéric Boulanger – Supélec (France), Dominique Potier (Thales R&T, France), Marius Minea - Institute e-Austria Timisoara (Romania), Julio Medina – University of Cantabria (Spain)

Overview: Existing component models and frameworks do not adequately support essential properties of real-time systems, such as heterogeneity, resources, behaviour, timing, and quality of service. Partners have been working towards a framework for component-based development of heterogeneous embedded systems, including the following approaches.

Work in Year 1

-- Design of Heterogeneous Systems --

A key characteristic of component-based embedded systems is heterogeneity of component models. This heterogeneity concerns different execution models (synchronous, asynchronous, vs. timed), communication models (synchronous vs. asynchronous), as well as different scheduling paradigms. The PARADES team has been a driving force in the development of the Metropolis (http://www.gigascale.org/metropolis) environment, which supports a variety of design notations and the concurrent management of different physical properties such as power, reliability, timing and cost. The design method is now being increasingly explored in the context of intelligent building, airplane engine, air conditioning systems and elevator design. The Metropolis environment supports the formal aspects of the design methodology.

VERIMAG has developed the Behavior, Interaction, Priority (BIP) framework for component-based modelling of heterogeneous real-time systems [Si05, BBS]. BIP integrates research results developed at VERIMAG over the past five years. It is characterized by the following:

It supports a component construction methodology based on the thesis that components are obtained as the superposition of three layers: (behaviour, interactions between transitions of the behaviour, scheduling policies for interactions). Layering implies a clear separation between behavior and structure (connectors and priority rules).

BIP has been successfully applied to define operational semantics for the Heterogeneous Rich Component model in the SPEEDS project.

TU Vienna has developed the foundations for an integrated architecture that facilitates the development of distributed real-time applications consisting of multiple heterogeneous subsystems with different criticality levels. A central issue is a framework for providing standardized, validated and certified services that can be reused in different applications.

-- Interfaces and Composability --

Several partners of the RTC cluster have been developing tools and techniques for specifying and reasoning about timing and resource properties of components and systems composed from components. These include the following.

- The MAST environment for schedulability modelling and analysis, which has been developed by the Univ. of Cantabria.
- The real-time calculus, developed by the team of ETHZ, which allows specifying components under less constraining assumptions, and represent many different kinds of properties (period, jitter, bursts) in a uniform way. A further advantage is that it supports separation of concerns, since computation resources are treated as first-class citizens along-side with functional and timing properties; the available computation resources are specified explicitly in a uniform representation.
- A more general technology for specifying and analyzing timing properties is offered by (variants of) timed automata. Several teams have developed tools for modelling and analysis of timed automata specification (UPPAAL by Uppsala and Aalborg, IF/Kronos by Verimag).
- An adaptation of automata-based techniques towards specifying components in terms of required and offered properties of their temporal behaviour is offered by the work on interface automata by the EPFL team and their collaborators. This work has also been extended to include quantitative timing properties as in timed automata in the work on timed interfaces.

Several partners have contributed to the development of component frameworks that can handle timing and resource properties. This has been done, e.g., in the on the Omega component model [DJPV05], Simpler component frameworks, which modestly extend existing mainstream techniques for design of real-time systems, include Rubus.

MdH, with contributions from Uppsala University, has developed SaveCCM, a component model and technology for real-time embedded systems with constrained resources. The technology aims for pratactibility of some properties and a transformation from the component model to an execution model optimised for resource usage (CPU or memory).

Work in Year 2

NB: The activity started during Year 2.

-- Design of Heterogeneous Systems --

The theory on *tag systems* was further developed by Benoît Caillaud and Dumitru Potop-Butucaru (VERIMAG, then INRIA, team Aoste), who developed a theory for the correct deployment of synchronous designs over globally asynchronous, locally synchronous (GALS) architectures. This work introduced the notion of weak endochrony, at a macro-step level, which extends to a synchronous setting the classical theory of Mazurkiewicz traces. A micro-step model for the representation of asynchronous implementations of synchronous specifications was introduced. The model covers classical implementations, where a notion of global synchronization was preserved by means of signaling, and globally asynchronous, locally synchronous (GALS) implementations where the global clock is removed. This model offers a more refined framework for reasoning about essential correctness properties of an implementation: the preservation of semantics and the absence of deadlocks. Stavros Tripakis and Paul Caspi of VERIMAG actively collaborated with INRIA and PARADES in developing techniques for heterogeneous systems modelling and in automatic code generation from high level synchronous models on several platforms, notably asynchronous preemptive ones.

The *BIP (Behavior, Interaction, Priority)* framework for modeling heterogeneous real-time components which integrates results obtained at *VERIMAG* over the past 5 years was implemented in a tool allowing the efficient execution of specifications. BIP is a central semantic-level formalism that is connected to several modeling formalisms and validation tools in the work of *Plaform for Component Modeling and Verification,* but is also an effort to enable integration of heterogeneous systems. A mapping from BIP to Think/Fractal is being implemented jointly with *FTR&D* for achieving code generation for BIP descriptions. Several industrial case studies have been modelled using BIP, including an Adaptive QoS controller for a video encoder, a planner for autonomous robots and we started to work on a model of sensor networks (together with FTR&D) for fine grained energy consumption analysis.

TU Vienna has worked on a next-generation embedded architecture for Systems-on-a-Chip (SoCs) that provides a predictable integrated execution environment for the component-based design of many different types of embedded applications (e.g., consumer, avionics, automotive, industrial). The architecture is inspired by the research priorities that have been identified in the ARTEMIS Strategic Research Agenda (SRA), such as composability, networking, robustness/security, diagnosis, resource management, and evolvability. The network interface will be based on the Time-Triggered Ethernet (TTE) protocol that supports the coexistence of hard real-time communication and standard Ethernet messages [KAGS05, OPK05]. The *OFFIS* team has developed an approach to design space exploration within the development of distributed embedded real-time systems. The mapping of software parts onto suitable hardware parts is a crucial issue of optimization towards efficient and inexpensive implementations. An extended SMT checker is used in a binary search scheme in order to achieve optimal allocations of tasks and messages to architectural elements.

-- Interfaces and Composability --

The work on developing the concept of *rich component models* into a mature framework for system design has been pursued within IP-SPEEDS by RTC partners INRIA, OFFIS, PARADES, and VERIMAG. A goal of SPEEDS is to provide an engineering environment enabling the creation, manipulation, and maintenance of rich component models and allowing system engineers to perform analysis, evaluate the maturity of the design and exchange design representations at different level of abstractions. Currently, the work is focussing on developing a meta-model for rich components. This includes defining a notion of component for which different *viewpoints* (functional, times, safety, etc) can be synchronized, and different viewpoints for different components can be formally composed. It will comply with existing or de-facto standards, including the Autosar real-time component model, UML 2.0 (in particular SysML profile). The work in SPEEDS also involves a new theory of *interfaces*, allowing for cross-viewpoint assume-guarantee reasoning. More precisely, a novel notion of contract has been defined for embedded systems that takes their multiple viewpoint nature into account. It was found that the way contracts should be composed for different viewpoints of a same component differs from the one used for different components. The fusion of contracts is a new operator that subsumes both cases.

Several lines of work have focussed on timing properties. Different techniques for specifying and analyzing timing properties, including the real-time calculus (developed at ETHZ), classical schedulability analysis, and timed-automata techniques (implemented, e.g., in Uppaal) have been compared in the the workshop “Distributed Embedded Systems” at the Lorentz Center in Leiden in Nov. 2005. A diploma project at Timisoara implemented a translation from a dedicated description language for multiprocessor tasks into Uppaal models using timed automata. Uppsala has developed a translation between the real-time calculus of ETHZ and timed automata formalism. This translation is currently being implemented in Uppaal. The *EPFL* team has developed an assume-guarantee interface algebra for real-time components. In this formalism a component implements a set of task sequences that share a resource. The

algebra defines compatibility and refinement relations on interfaces. The algebra thus formalizes an interface-based design methodology that supports both the incremental addition of new components and the independent stepwise refinement of existing components. The flexibility and efficiency of the framework has been demonstrated through simulation experiments.

Integration of techniques from schedulability analysis into component-based design methods are further developed by *Cantabria* and *Thales* in the FRESCOR project: Framework for Real-time Embedded Systems based on COntRacts (www.frescor.org, IST-034026), which aims to produce a framework for handling timing requirements with a focus on reconfigurable architectures. Within the context of the SAVE Swedish national project, the Uppsala and Mälardalen teams are developing *SaveCCM* (the SaveComp component model).

EPFL and PARADES have collaborated to adapt techniques for specifying component interfaces for the development of a structured coordination language for specifying the interaction of real-time tasks. Task communication happens through shared variables called communicators, which can be read and written only at specified time instances. Sensors and actuators are special kinds of communicators. The read and write times of communicators determine the release times and deadlines of tasks. Tasks may also depend on each other, be refined into sets of tasks, and be changed through mode switches. The language is a hierarchical extension of Giotto, and has been inspired by and used in the automotive domain.

Dortmund and Uppsala have collaborated to develop and implement automata learning techniques for automatically deriving behavioural models of components from legacy code or observations of system behavior. Part of the work concerns extending these techniques to derive timed models.

-- Industrial Liaison --

The forums organized in the framework of this activity are an important contribution to the interaction between industry and academia in the considered sector. The meeting *Meeting Beyond AUTOSAR* was held on March 23rd - 24th, 2006 in Innsbruck, Austria. There were 52 registered participants, among which 15 from industry. The agenda of the meeting, as well as the detailed minutes and slides can be found at

http://www.artist-embedded.org/artist/-ARTIST2-Workshop-Beyond-AutoSar-.html

Work in Year 3

-- Design of Heterogeneous Systems --

An algebraic framework for BIP (VERIMAG)

We worked for an algebraic formalization for the BIP framework [BS07]. The main difference with existing process algebras is the use of operators for composing connectors describing interactions and priority. We provided an algebraic formalisation of connectors in BIP. These are used to structure interactions in a component based system. A connector relates a set of typed ports. Types are used to describe different modes of synchronisation: rendezvous and broadcast, in particular. We used the system construction space *Behaviour* × *Interaction* × *Priority* to study relations between different classes of models. We studied in particular, characterizations of existing models of computation as regions of this space and relations between these regions. Furthermore, different subclasses of models e.g., untimed/timed, asynchronous/ synchronous, event-triggered/data-triggered, can be unified through transformations in the construction space.

Designing a timed BIP component model (INRIA and VERIMAG)

Verimag and INRIA have collaborated by merging their component models to design a timed BIP component model. This model will be integrated in the platform under construction by the Platform activity of the CBD cluster [SPB07].

A formal approach for modelling heterogeneous systems (CEA LIST)
In order to transfer the research on modeling of heterogeneous systems into the standardization domain, work started to build a specialisation of a standard modelling language (UML) to describe heterogeneous computation and communication models founded on a mathematical basis. This lead to create a common research action (THeSys: Tackling Heterogeneous Systems – www.thesys.eu.org) of the research cluster Digiteo Labs (www.digiteo-labs.org) with Suppelec, ARTIST 2 associated member. A first report on the state of the art was produced for Usine Logicielle and a first instantiation of a dedicated UML profile was provided.

Architecture for Heterogeneous Systems (TU Vienna)
The important aspects on error containment and diagnosis within heterogeneous distributed systems were addressed within [OKSH07] and [EOHPK07], as a continuation of the work on diagnosis that was started in year 1 in the HRT cluster. The proposed architecture enables the integration of mixed-criticality subsystems (cf. [EOHKS07]) within a distributed system and even within a single chip. Other problems addressed by the Vienna team included the following

- To facilitate the modeling and formal verification of distributed heterogeneous systems which are designed according to the time-triggered paradigm, The Periodic Finite-State Machines (PFSM) [KEHO07] were developed as extensions of the basic Finite State Machines (FSM) model, based on the concept of a sparse time base.
- [OH06] presented a solution for the model-based design of virtual networks in distributed heterogeneous networks enabling faster development time and avoiding design faults. This work was extended to an overall model-based development process of integrated computer systems based on the DECOS architecture in [HO07].
- [SOE07] attacked the problem of interfacing heterogeneous distributed applications to Hardware-in-the-Loop (HIL) simulators and presents a solution based on an interface at the sensor/actuator level.

-- Interfaces and Composability --

Several interacting lines of work hwere performed in the context of efforts where component models for embedded system design are developed. The work on *Rich Component Model* in SPEEDS targets both heterogeneous and component-based systems. Other efforts (described subsequently) are more focussed on timing and resource problems in component based design.

-- Meta model for Heterogeneous Rich Components (INRIA, OFFIS, PARADES, VERIMAG) --
Within the IP SPEEDS, the work on developing the *Rich Component Model* paradigm was focussing on the development of a metamodel, called HRC (Heterogeneous Rich Components), which now forms the foundation for the component based construction of complete virtual system models. Its main objectives are: 1) to define a semantic-based meta-model used by all involved tools, 2) to develop a framework for multiple viewpoint (functional and non-functional) component engineering, 3) to enable full-scale reuse of components, 4) to offer from COTS modelling tools, access to meta-model compliant components and, 5) to assess early project risks at subsystem level to secure concurrent design processes.

During the first year of the project, a first version of this meta-model was defined. The main features of HRC are:

- *Design by contract* paradigm: attached to a component, contracts express constraints on assumed behaviour of the environment (assumption) and expected behaviour of the component (promise).
- *Organization in viewpoints: f*ollowing the principle of separation of concerns, different aspects are organized into viewpoints, each of which collects a part of the component's dynamics constraints from some perspective and can be used to filter the component's characteristics w.r.t. that view.
- *Uniform concepts across all layers:* different *layers* may be identified for expressing different architectural abstractions of an embedded system. Examples of layers are the functional layer representing the functionality of the system and the platform layer that together with the functional layer abstracts the system as a network of buses and ECUs (containing tasks and threads)..
- *Rich connectors*: in addition to SysML-like connectors expressing data or event flow with a unique predefined initiator, HRC contains more powerful connectors whose activation depends on the agreement between at least a subset of the connected components.

-- Validation and design space exploration--

Different specific validation techniques were developed or adapted for HRC models. We mention as examples, efficient deadlock analysis using the structure provided by rich connectors, simulation using BIP (Verimag) or Metropolis (parades), Hybrid analysis using Ariadne (Parades), and specific methods for timing or safety analysis (OFFIS). They are reported in more details in the platform deliverable. OFFIS and PARADES collaborated on design space exploration based on HRC. Here, the deployment of executable components and communication links determines the extra-functional properties, such as timing. Finding a cost efficient and requirement preserving deployment is subject of optimization. The deployment synthesis OFFIS developed (RTSat) provides the capability of finding optimal deployments among the solution space for a given architecture, while preserving extra-functional requirements on real-time, memory, etc, which were shown to be fulfilled at the specification level [MH06].

-- The SaveComp component model --

Another effort to develop a model for component based development is *SaveCCM* (the SaveComp component model), developed by the Mälardalen and Uppsala teams [ÅCF+07]. SaveCCM is based on a control-flow (pipes-and-filters) interaction model, combined with additional support for domain specific key functionality. Timing properties of a system of components can be analyzed using fixed-priority analysis techniques, using e.g., the MAST schedulability modeling and analysis environment developed by the Univ. of Cantabria. The *SaveCCM* component model has been employed in industrial case studies, e.g., at CC Systems, where a component-based repository is being built.

-- Deployment of LightWeigth CCM components within a Flexible scheduling framework. --

In the context of the effort to combine real-time implementation technology and contract technology to build techniques for component-based design, in the context of the FRESCOR project, University of Cantabria and Thales used a specialization of the Deployment and Configuration OMG standard to define an approach for the deployment of MicroCCM components. The initial design [LPDM07] was made by Patricia López from Cantabria and allows the generation of the analysis models from the same description.

-- Hierarchical Coordination Language for Interacting Real-Time Tasks --

As another concrete technology for component based development, EPFL and PARADES designed and implemented a new programming language called Hierarchical Timing Language (HTL) for hard real-time systems. HTL is a hierarchical version of Giotto. Critical timing constraints are specified within the language, and ensured by the compiler. As a case study, a distributed HTL implementation of an automotive steer-by-wire controller was implemented [GHIKS06].

-- Platform implementation technology for timed components --

The work on a transformation chain for timed components was extended to allow assembly and automatic mapping onto the Giotto framework. The tool is able to accept assemblies of timed components, check the assemblies for compliance with timed logic properties and generate a set of monitor for execution on these assemblies on the Giotto infrastructure from EPFL. Moreover, the INRIA team also designed a special version of a Java machine able to run on the Lego Mindstorm platform (a tiny, low cost commercial platform for building robots). Mindstorm software is monitored in situ using automatically generated monitors. This joint work merges the advantages of BIP components on structuring and continuous time management [SBD06].

-- Scalable Specification and analysis of timing properties --

The work conducted on developing techniques for analysis of timing and resource properties, which are more precise and more scalable than existing ones was continued in Y3 with an implementation of translations between the real-time calculus, developed at ETHZ, and timed automata formalisms in the context of the Times tool (http://www.timestool.com). A prototype tool (named CATS) for compositional timing and performance analysis was developed.

Within software engineering for embedded systems generic reusable software components must often be discarded in favor of using resource optimized solutions. In cooperation with the Swedish company CC Systems, MdH has developed a model that enables the utilization of component-based principles even for embedded systems with high optimization demands. The model supports the creation of component variants optimized for different scenarios, through the introduction of an entrance preparation step and an ending verification step into the component design process. These activities are proposed to be supported by tools working on metadata associated with components, where the metadata can be automatically retrieved from many development tools [ÅFSC07]

-- A Model for Reuse and Optimization of Embedded Software Components --

In cooperation with the Swedish company CC Systems, MdH developed a model that supports the creation of component variants optimized for different scenarios, through the introduction of an entrance preparation step and an ending verification step into the component design process.

-- Adapter synthesis for real-time components --

An approach for overcoming compatibility problems in composition of available components, was developed by INRIA and L'Aquila University. An automated method was devised to build correct-by-construction adapters, to be inserted between components such that all inconsistencies are solved. This is possible thanks to the controllability of some input and output actions. The method uses a Petri Nets modelling and a specific controlled coverability graph generation algorithm. It is implemented inside a tool suite [TFGG07].

-- Algorithms for Interface Synthesis (EPFL, Uppsala, and Dortmund) --

With the goal to extend the available repertoire of techniques for generating component models, Dortmund and Uppsala have beencollaborating to develop automata learning techniques (aka regular inference) for automatically deriving behavioural models of components from observations of system behavior. Such techniques can be useful to generate models of components for which no source code is available, e.g., libraries, hardware components. Dortmund has developed *LearnLib* [BRS06], a library for automata learning, with a flexible modular structure that can be configured to exploit specific properties of applications. During Y3 of ARTIST2, the collaboration has been motivated by the goal of using LearnLib to generate a model of an industrial protocol developed by an industrial partner of Uppsala (Mobile Arts AB). One difficulty in this protocol is that messages contain identifiers of connections, etc. from a potentially infinite domain. The main achievement during Y3 has been to extend automata learning techniques to a class of infinite-state systems that can handle this situation [BJR]. Another line of work concerns extending automata learning techniques to generate models of timed systems, in the form of timed automata [GJP06].

EPFL compared and evaluated three different algorithms for automatically extracting temporal interfaces from code: (1) a game algorithm that computes the interface as a representation of the most general environment strategy to avoid a safety violation; (2) a learning algorithm that repeatedly queries the program to construct the minimal interface automaton; and (3) a CEGAR algorithm that iteratively refines an abstract interface hypothesis by adding relevant program variables. For each of the three algorithms a family of components was provided on which that algorithm outperforms the two alternatives. On the practical side, the three algorithms were evaluated experimentally on a variety of component libraries [BHS07].

-- Industrial Liaison --

Organization of Workshops on Industrial Topics
The workshop “Beyond AUTOSAR” held in the Year 2 period gathered key industry players from AUTOSAR and key scientists to discuss fundamental issues for embedded automotive systems design. Werner Damm presented the results of the workshop in a keynote lecture at the EMSOFT Conference 2006 in Seoul and at a workshop organized by GM on the Future of Automotive Software Development in Bengalore (January 2007). The documentation for the findings of the workshop are at the “proceedings site” http://www.artist-embedded.org/artist/-ARTIST2-Workshop-Beyond-AutoSar-.html

PARADES is an industrial research consortium. Its partners (Cadence and ST) are constantly made aware of the technical advances pursued by the PARADES team. The interaction with people in the companies is at least weekly. ST has a strong interaction on fault tolerant architectures and fault analysis and uses PARADES expertise to interact with system customers such as Bosch and Nippon Denso. PARADES is also in contact with Freescale via the Joint Development Group with ST. Cadence relies on PARADES expertise for system-level design methodologies and tools. PARADES has interacted with Pirelli in a project involving intelligent tires for stability control in cars. PARADES has had significant interaction with United Technology Corporation (UTC), a large multi-national conglomerate, on sponsored research for embedded system architecture and design methodologies for OTIS Elevators, Carrier air conditioning systems and Chubb Securite’, a large division in charge of safety and security systems for buildings and large structures such as hospitals. In addition, PARADES has had collaboration with General Motors on research strategies and directions.

Establishment of SafeTRANS
During Y3, OFFIS was instrumental in creating SafeTRANS (http://www.safetrans-de.org), a non-profit organisation combining the expertise of German key industrial and academic players in the area of processes and methods for the development of safety critical embedded systems

in the transportation domain. Building on OFFIS' strong industrial cooperation network and using experience gained from numerous activities in shaping European R&D roadmaps, SafeTRANS founding members are Airbus Germany, Bosch, Continental, DaimlerChrysler, Siemens VDO and Transportation Systems, OFFIS, DLR and the Carl von Ossietzky University of Oldenburg. SafeTRANS' mission is to to maintain the current high safety levels of transportation systems in spite of growing traffic density, and in spite of an exponential growth in Embedded Systems complexity, through model based development and analysis of safety-critical Embedded Systems enabling a holistic system analysis.

Together with two french Pôle de Compétitivités Aerospace Valley (http://www.aerospace-valley.com) and System@tic (http://www.systematic-paris-region.org), SafeTRANS formed EICOSE, the European Institute for COmplex and Safety Critical Embedded Systems Engineering[1]. Through the participating competence centres, EICOSE clusters major industrial and academic organisations in the area of embedded systems in the transportation domain, namely Airbus, Alcatel Space, Alstom, Altis, Astrium, Bosch, CEA, Cegelec, CNES, CNRS, Continental, CS communication et Systèmes, DaimlerChrysler, Dassault-Aviation, Dassault Systems, DLR, EADS ST, EDF, ENSC, Ecole Polytechnique, France Telecom, IERSET, INRIA, IRC SCS, LAAS, Latécoère, Motorola, OFFIS, ONERA, RATP, Renault, SiemensVDO, SiemensTransportation, SNCF, SNECMA, Sogerm, Thales, University of Oldenburg, Valeo, Visteon, and many others. EICOSE has been selected the first ARTEMIS Innovation cluster, paving the way for EICOSE to participate in shaping those parts of the ARTEMIS Strategic Research Agenda dealing with the transportation domain, thus directly influencing calls in the forthcoming ARTEMIS JU. EICOSE has identified a priority list of research items from an industrial point of view, which lead to the formation of proposals of three so-called subprograms which were adopted by Artemis in its Artemis Multi-Annual Strategy Plan, and included in the call for proposals for the 1st call of the Artemis Joint Undertaking.

Final Results

-- Design of Heterogeneous Systems --

Formalisms for modelling Heterogeneous Systems: The partners' different approaches to modelling heterogeneous systems have been further developed within the context of ARTIST2. Results achieved include

- Further work on Tag Systems (INRIA, Parades, and Verimag):
  The theory of Tag systems has been further developed [BCCCS08]. Tag systems are models of systems where data are enriched with tags, supporting a flexible and parameterizable notion of time. This approach supports heterogeneity by providing a mathematical basis for composing subsystems with different Models of Computation and Communication (MoCCs). Theorems have been provided that give conditions for the original semantics to be preserved at deployment phase, when a "less synchronous" architecture is used. This is work along lines similar to the efforts of Edward Lee regarding Ptolemy II. Ongoing work now includes participants from Cadence Berkeley. It consists in developing a functional version of Tag systems theory that strongly relies on Kahn Process Network techniques. In particular we were able to show that directors (in the sense of Ptolemy II) are not needed in order to coordinate different MoCCs in this family.
- A notion of expressivity for composition formalisms (Verimag):
  In Year3 and 4, the conceptual work on BIP has focussed on the development of an

[1] http://www.artemis-office.org/DotNetNuke/Activities/EICOSE/tabid/123/Default.aspx

algebraic theory [BS07a, BS07b]. In the final year, Verimag has also proposed a new notion of expressiveness appropriate for formalisms that express composition of components. It compares component frameworks with respect to the ability to achieve new behaviours from a given set of component behaviours. The work proposes an SOS-style definition of glues, where operators are characterized as sets of SOS-rules, specifying the transition relation of composite components from the transition relations of their constituents. We provide expressiveness results for the glues used in BIP and for process algebras such as CCS, CSP and SCCS. We show that the glues used in CCS, CSP, and SCCS are less expressive than the general SOS glue, but that BIP has the same expressiveness as SOS-style rules, which means that when restricting to memoryless composition or glue operators, BIP has maximal expressivity [BS08]. This is an indication that the concepts of BIP are as general as one may need.

- Description of models of computation and of models of execution (CEA, Supélec):
  The formal LEM language (PC-xUML) [CGMOBHM07] for modelling models of computation is now implemented as a UML profile available in the Papyrus modeler. The ModHel'X framework [BH08], which provides a generic meta model for describing heterogeneous systems, a generic execution engine for simulating heterogeneous models, and a language for describing models of execution and their interactions is also available in a preliminary version at http://wwwdi.supelec.fr/logiciels/modhelx/. We consider LEM as a language for specifying models of computation (rules of combination of behaviours), and ModHel'X as a framework for describing models of execution (algorithms for combining behaviours), which lead us to study the conformance of a model of execution to a model of computation in a way similar to the conformance of an implementation to a specification. These works are conducted in the context of the TheSys research group http://www.thesys.eu.org .

New Models of Computation (INRIA):
Models of Computation was the focus of an ARTIST2 workshop in Zürich in Nov. 2006, which inspired further workon the topic. INRIA has proposed a novel Model of Computation (MoC): Kahn-extended Event Graph (KEG) which add "static control" (control known at compilation time) in the MoC Marked Graphs (MG). INRIA has also defined a process of *expansion*, which finds the parallelism in a model and transformes it to an “expanded” (more parallel) one. The approach has been illustrated on a simple C algorithm (a Sobel filter) [BCFMS08]. INRIA is also building a new version of their tool K-PASSA, which finds static schedules of system descriptions, in order to add the following MoCs: KEG, Synchronous Data Flow (SDF) and Latency-Insensitive Design (LID). A previous release was implementing Marked Graphs (MG) and a specific optimization called "equalization". In the context of the French regional CIM PACA collaborative center some of the results have been demonstrated to industrial partners (such as Texas Instruments, ST Microelectronics, NXP, Synopsys, and smaller French SMEs). INRIA is currently investigating the relevance of the KEG models and their associated static schedules for the design and optimization of Networks-on-Chip traffic. http://ralyx.inria.fr/2006/Raweb/aoste/uid27.html

-- Distributed Implementation of non-distributed specifications --

The problem of realizing a distributed implementation of a non-distributed system description is a challenging problem, which has received attention by ARTIST partners, e.g., in the work on GALS (Globally Asynchronous, Locally Synchronous) systems. There are still many unsolved problems in establishing the foundations for distributed implementations. Progress at the end of ARTIST2 includes the following.

- A distributed semantics for BIP (Verimag):
  The operational semantics of the BIP language has originally been defined in such a way that interactions – defined by a data exchange between a set of components that is

followed by local steps of individual components – are executed atomically. This means that the decision about the set of possible next steps is posed only in states in which all components are in a stable state. This semantics has been implemented previously in the BIP engine. In order to support distributed implementation, Verimag proposes two alternative semantics allowing a pipelined execution of atomic steps. In both semantics, additional intermediate partial states are introduced by cutting each transition of an individual component into two steps such that in the new intermediate state a component is not ready for any communication – meaning that in such a state the global state is only partly defined. The first semantics ignores the distinction between partially defined and global states and computes the set of enabled transitions in a state using the information on components in a defined state only. The second semantics implements interactions in the partial state model by using message passing primitives. The main result of the work consists of conditions for which the models are observationally equivalent. We study performance trade-offs and provide experimental results illustrating the application of the theory on a prototype implementation [BBBS08].

- A new track on Loosely Time-Triggered Architectures (LTTA) (INRIA, Parades, and Verimag):
  LTTA aims at relaxing the strict synchrony constraint of Kopetz' TTA by allowing local clocks of computing and communication units not to be synchronized. LTTA architectures are widely used (even more than strict TTA) in industrial control such as flight control, nuclear plant monitoring, railway control... and most of their programming uses synchronous formalisms (Simulink, SCADE/Lustre and similar control-based formalisms). We have studied several variants and have shown how specification semantics can be preserved [BCNPST07] [TPBSCN08] [CB08]. Verimag and INRIA have further developed initial work by Verimag on the study of Airbus system architecture for low level flight control. They have come up with the systematic idea of replacing token based mechanisms by the use of purely local counters with no additional link. Each unit maintains a local counter based on its own independent clock. This local counter controls the right to acquire new input data from the communication media, perform computation steps, and write output data to the communication media. This approach is entirely local. Pros and Cons of this approach are:
  - Pros: no back-pressure, no additional communication link, no blocking communication; this simplifies the design of fault-tolerance and degraded modes.
  - Cons: uses boundedness assumptions on the relative drift between local clocks (the management of the local counters depends on these bounds). This means a high cost when re-design is needed.
- Implementing synchronous models on asynchronous architectures (PARADES, INRIA, SSA, and Verimag):
  This item addresses the same question as the previous one by moving from loosely synchronised to fully asynchronous architectures and share in common some solutions. In collaboration with Cadence Berkeley Research Labs, and UCBerkely, a theory for the design of communication architectures has been developed, that would guarantee the same property as a synchronous architecture but would be implemented on an asynchronous one [MBFS07,TPBSCN08]. PARADES, INRIA, and Verimag, jointly with UC-Berkeley and Cadence Berkeley, have developed approach (a), resulting in publications [B&al07] and [T&al08]. The approach assumes a single-clocked synchronous specification – single-clocked is not really a restriction in this context as it can be relaxed by using the extra symbol *nil* meaning the absence of a certain data at a given reaction. It is known that such specifications can be seen as a Kahn Process

Network with bounded buffers. This observation has been the basis for the development of so-called *elastic circuits* by Cortadella et al. and *latency insensitive designs with back-pressure* by Carloni and Sangiovanni-Vincentelli in the area of circuit design. These are circuits with token based mechanisms. Controlling buffer overflow is achieved by implementing backward tokens controlling the permission to write in buffers – hence the term of *back-pressure.* This idea has been adapted to our case where neither writing nor reading can be blocking, see the figure above. The idea is to replace blocking by skipping. Performance of such architectures is classically studied by means of Marked Graphs, a simple form of Petri nets where Max-Plus algebra applies. Pros and Cons of this approach are:

- Pros: no assumption on local clocks; very adaptive, scales up easily to complex systems; easy upgrade.
- Cons: need for back-pressure, which results in additional links, resulting in additional requests for fault tolerance mechanisms.

- Reliability of distributed implementations (EPFL, PARADES):
  EPFL and PARADES have designed and implemented a hierarchical version of the programming language Giotto, called Hierarchical Timing Language (HTL) for hard real-time systems. (see Section 2.3 on Y3 results). In Y4, the HTL framework has been extended to handle reliability. More precisely, EPFL, PARADES and UCB proposed an abstract notion of logical reliability for real-time program tasks that interact through periodically up-dated program variables. With each program variable is associated a *logical* (or *long-term*) *reliability constraint* (LRC), a real number between 0 and 1. If the LRC is 0.9, this means that in the long run, at least a 0.9 fraction of all periodic writes to this communicator are required to be valid values. The mapping of tasks to hosts must ensure the LRCs of all program variables. For this purpose, if hosts fail, it may be necessary that a task be replicated on several hosts. To check if an implementation satisfies all LRCs, the singular (or short-term) reliability guarantee (SRG) of updating a program variable with a valid value must be known. The SRG is again a real number between 0 and 1; for example, an SRG of 0.8 means that the probability that a host fails during the execution of a task invocation is 0.2. The SRG is a property of the architecture, just as WCETs are architectural properties. To achieve LRCs of 0.9 with hosts that guarantee only SRGs of 0.8, all tasks that write to communicators (with LRC 0.9) need to be replicated on two hosts. The HTL compiler has been extended to perform also a reliability analysis and to generate distributed code that satisfies the requirements [CGH+08] [PCS08].

- New Heuristics in Scheduling for Reliability (INRIA):
  As a contribution to scheduling for distributed reliable real-time systems, INRIA has proposed a new framework for the (length, reliability) bicriteria static multiprocessor scheduling problem. The first criterion remains the static schedule's length: this is crucial to assess the system's real-time property. For the second criterion, we consider the global system failure rate, seen as if the whole system were a single task scheduled onto a single processor, instead of the usual reliability, because it does not depend on the schedule length like the reliability does (due to its computation in the classical exponential distribution model). Therefore, we control better the replication factor of each individual task of the dependency task graph given as a specification, with respect to the desired failure rate. Compared to the other bicriteria (length, reliability) scheduling algorithms found in the literature, the algorithm we present here is the first able to improve significantly the reliability, by several orders of magnitude, making it suitable to safety critical systems [GK08].

Apart from conceptual contributions to the problem of distributed implementation, there is also a need for design principles and architectures for concretely realizing distributed implementations, typically on NoCs. Most VLSI circuits can be considered distributed systems.

Since their components are designed independently, the assembly step is often a challenging problem that requires the design of communication interfaces to match different protocols and data parallelism, and the routing of global interconnect wires to meet the constraints imposed by the target clock period. The debate between those who favor standard bus architectures or variations thereof and those who advocate the adoption of NoC approaches ranging from constrained architectures to custom ones is vibrant.

- Design of Communication Architectures (PARADES):
  UC Berkely, Columbia and PARADES, developed a common framework, COSI, for the synthesis of communication for distributed systems, including chips as well as buildings.The proposed approach embedded in COSI does not take sides even though the NoC approach has undisputable fundamental merits that may make it successful in the long run. Instead, COSI proposed a general methodology for the design of on-chip communication that can explore a large number of alternatives including as special cases NoCs, bus architectures and hybrid ones. Thanks to its generality the approach can be used to build a framework where different constrained solutions are compared using a number of evaluation factors. Models for functionality, cost, and performance of each element are captured in the library together with their composition rules. A mathematical framework was developed to model communication at different levels of abstraction from the point-to-point input specification to the library elements and the final implementation. The code is publicly available: http://embedded.eecs.berkeley.edu/cosi/Home.html

- The Time-Triggered System-on-a-Chip (TTSoC) architecture (Vienna):
  is a novel system architecture that enables the realization of mixed-criticality systems using SoCs. It represents the culmination of several years of effort by the TU Vienna team (see also previous results from previous years). The integration of subsystems with different criticality enables massive cost reduction by reducing the overall number of devices and networks (e.g., ECUs in car). To accomplish this goal, the TTSoC architecture offers inherent fault isolation mechanisms that prevent any unintended interference between application subsystems of different criticality. Vienna has demonstrated these capabilities using an automotive example with a safety-critical control subsystem and a multimedia subsystem. In the demo application, it is ensured by construction that any design fault in the multimedia subsystem cannot have any adverse effect on the safety-critical control subsystem. The central element of the presented System-on-Chip (SoC) architecture is a time-triggered Network-on-a-Chip (NoC) that interconnects multiple, possibly heterogeneous IP blocks called micro components. The SoC introduces a trusted subsystem, which ensures that a fault (e.g., a software fault) within the host of a micro component cannot lead to a violation of the micro component's temporal interface specification in a way that the communication between other micro components would be disrupted. For this reason, the trusted subsystem prevents a faulty micro component from sending messages during the sending slots of any other micro component. Furthermore, the time-triggered SoC architecture supports integrated resource management, and failure detection, masking, and encapsulation using Triple Modular Redundancy (TMR). In summary, The SoC architecture ensures *composability:* that upon the incremental integration of micro components, the prior services of the already existing micro components are not invalidated by the new micro components [KEHOP08,OEHK08,OFEH08,OKS08].

-- Interfaces and Composability --

Interface theories with component reuse (EPFL):
Interface theories have been proposed to support incremental design and independent implementability. Incremental design means that the compatibility checking of interfaces can

proceed for partial system descriptions, without knowing the interfaces of all components. Independent implementability means that compatible interfaces can be refined separately, maintaining compatibility. General theories, which do not focus on a specific formalism for specifying interfaces but rather on what such formalisms can do, for interface-based design have been proposed, e.g., by EPFL (see report for Year 2). We have now shown that these interface theories provide no formal support for component reuse, meaning that the same component cannot be used to implement several different interfaces in a design. We therefore added a new operation to interface theories in order to support such reuse. For example, different interfaces for the same component may refer to different aspects such as functionality, timing, and power consumption. We gave both stateless and stateful examples for interface theories with component reuse. To illustrate component reuse in interface-based design, we showed how the stateful theory provides a natural framework for specifying and refining PCI bus clients [DHJP08].

Reasoning about systems of components.
The problem of analyzing or verifying a system of components has continued to receive attention in several contexts developed by ARTIST2 members.

- Contract-based verification for the Heterogeneous Rich Component (HRC) Model (INRIA, Parades, and Verimag):
  In Y3, the SPEEDS project defined a meta-model for representing component systems which includes a notion of contract that can be attached to components (see Section 2.3). Work has progressed during Y4 [BCP07] [B08] [CMMSS08] [M etal08] [JM08] [Met08] [DJMNKSV08] [DM07]. We have defined a satisfaction relation between an implementation and a contract and a notion of refinement between contracts. Initially, these relationships have simply been defined in terms of inclusion between trace sets. We have also developed a general framework for contract-based reasoning, which handles (multipartner) rendez-vous – as in the HRC framework – as well as many different languages for describing behaviours notions of refinement under context [QG08]. We use BIP to represent contracts. First, we define a general notion of a component framework defined by (1) a set of composition operators (2) a family of possible behaviours, (3) a mapping from *n*-ary composition operators into *n*-ary composition operators on behaviours, and finally (4) a notion of refinement under context between behaviours. We provide then several sufficient conditions which can be used to prove refinement in any framework that satisfies a certain property (“allowing circular reasoning”). We consider two particular instances of such a framework: the first are I/O automata with the usual composition operator and a notion of refinement based on trace inclusion. The second one considers behaviours defined by modal transition systems, allows all composition operators definable in BIP and a notion of refinement under context obtained from the usual notion of simulation between modal transition systems. At the review of the Integrated Project Speeds, the project demonstrated the capability of integrating models from multiple commercial of the shelf modelling tools based on the SPEEDS HRC meta-model, as well as running various analysis methods on the HRC representation of such integrated models, including checks for refinement, design space exploration, contract based safety analysis and real-time analysis, and hosted simulation.
- Compositional deadlock detection/verification (Verimag):
  Verimag has continued its work on deadlock detection/verification and its implementation in the DeadlockFinder tool by combining structural analysis for component behaviours with structural analysis of connectors [BBSN08].
- Integrating Modular Performance Analysis and Timed Automata Techniques (ETH, Uppsala)
  During Year 2-4, Uppsala and ETH have been conducting work on combining the

advantages of Modular Performance Analysis (MPA), which is based on the real-time calculus, and techniques based on timed automata. A prototype tool (named CATS) for compositional timing and performance analysis has been developed further during year 4. In CATS, a component can be characterized by equations over timed streams. The CATS tool is available at http://www.timestool.com/cats, and integrated in the Eclipse platform. During Y4, attention has also been given to handling cyclic dependencies in component-based real-time systems, which have not been well-understood in the context of modular performance analysis (MPA). To address this problem, a solid semantic foundation for MPA should be developed, linking operational behavior with the stram-based approach in MPA. ETH and Uppsala has developed a general operational semantics underlying the Real-Time Calculus, and used it to show that the behavior of systems with cyclic dependencies can be analyzed by fixpoint iterations. The work also characterizes conditions under which such iterations give safe results, and also show how precise the results can be [JPTY08].

Component Models:
Existing approaches to enhancement of existing component models have been continued.

- **Unified component model for embedded middleware (CEA, INRIA, THALES, STMicroelectronics):** *The* First version of a unified common meta-model for CCM, UML, MARTE, Fractal and OASIS component models has been proposed (www.flex-eware.org).
- **ProCom component model for distributed embedded systems (MdH, Uppsala):** The work on SaveCCM by MdH and Uppala has continued with a component model suitable for distributed embedded systems, resulting in the ProCom component model. The ProCom component model is developed i) for scalable design of small or large embedded systems, ii) for integration of different models for prediction and analysis of components and system properties, and iii) to allow resource-efficient realizations at run-time. The desired characteristics have been obtained by designing a two-layered component model where the lower layer strictly defines the execution semantics and enables efficient timing and resource analysis, while the top level enables a variety of component designs and styles of communication. Together with ProCOM the Progress Integrated development Environemnt (Progress-IDE) is being developed, including, e.g., the UppaalPort analysis tool [HMP07]. In addition to the component model, MdH has developed techniques supporting the design and development, integrated with the component model. Some of them are: i) Context aware execution-time estimation ii) Stack-sharing in component-based systems, iii) Advanced flow analyis iv) Partial order verification v) Software component-based development process vi) Formalization and automation of component selection. [SVBCC08] http://www.mrtc.mdh.se/progress/

Deployment of LightWeight CCM components within a Flexible scheduling framework (Thales, Univ.Cantabria):
In the context of the effort to combine real-time implementation technology and contract technology to build techniques for component-based design, in the FRESCOR project, University of Cantabria and Thales have continued their work on using a specialization of the Deployment and Configuration OMG standard to define an approach for the deployment of MicroCCM components [GHC+08] . A model based technology aiming at Ada implementation has been proposed [LDPM08]. The concept of interface in Ada 2005 significantly facilitates its usage as the basis for a software components technology. This technology, taking benefit of the resources that Ada offers for real-time systems development, is suitable for component-based real-time applications that run on embedded platforms with limited resources. The proposed technology uses the specification of components and the framework defined in the LwCCM standard, modifying it with some key features that make the temporal behaviour of the applications executed on it, predictable, and analysable with schedulability analysis tools. The dependency on CORBA is replaced by specialized communication components called

connectors. The threads required by the components are created and managed by the environment, and interception mechanisms are placed to control their scheduling parameters in a per-transaction basis. This effort aims to proposing a new IDL to Ada mapping, a prospective standard of the OMG. http://www.ctr.unican.es/publications/plm-jmd-ppm-jlm-2008a.html.

Synthesis of Glue and Controllers from Specifications (EPFL, INRIA):
This problem has been addressed by several lines of work. Controller synthesis problems are naturally formalized by games where one player (the program) has to entail some objective (the specification) no matter how the other players (other programs and external environment) behave. The winning strategy in such a game is a model of the controller to synthesize.

- EPFL has contributed several results on solving games, among them on timed games [CHP08].
- INRIA has proposed a schema of integrating Discrete Controller Synthesis (DCS) techniques into the modular compilation of an extended synchronous language. In this extended language, modularity is expressed by nodes, representing components associated with modular synthesis objectives; we can then obtain, by application of DCS tools on these components, some synchronous controllers controlling parts of programs. In this framework, we have implemented a translation schema of a subset of the Lucid Synchrone language into dynamic systems, for further application of Sigali, as DCS tool. Future work will consist in applying decentralized control methods, together with modular distribution of synchronous programs, in order to obtain automatically, from an annotated synchronous program, a distributed controlled system.
- INRIA has, together with Univ. of Auckland , developed a technique for synthesizing glue logic, termed as a converter, so that the parallel composition of the components and the converter also satisfies some desired specification. A converter is responsible for bridging different kinds of mismatches such as control, data, and clock mismatches. Mismatches are usually removed by the converter (similar to controllers in supervisory control of discrete event systems) by disabling undesirable paths in the protocol composition. We have generalized this convertibility verification problem, by using a new refinement called specification enforcing refinement (SER) between a protocol composition and a desired specification. The existence of such a refinement is shown to be a necessary and sufficient condition for the existence of a suitable converter. We have also proposed an approach to automatically synthesize a converter if a SER refinement relation exists.

Generating component models from observations of behaviour (Dortmund, Uppsala): Dortmund and Uppsala have been collaborating to develop automata learning techniques (aka regular inference) for automatically deriving behavioural models of components from observations of system behavior. They are intended to be used in situations where models or specifications are not available a priori, and where static source code analysis is not feasible.Potential uses of such models is in regression testing, as a guide for model based test suite generation, in generating models of systems and of component environments for various purposes. The work has resulted in the tool *LearnLib* [BRS06], mainly developed by Dortmund, which is a library for automata learning, with a flexible modular structure that can be configured to exploit specific properties of applications, in order to make automata learning scalable to realistic settings. During Year 4, standard automata learning techniques have been extended to a class of infinite-state systems that can handle this situation [BJR08]. Furthermore, work on developing techniques that make modelling of industrial protocols practical are in completion [BJ]. Finally, the work on.the work extending automata learning techniques to generate models of timed systems has been thoroughly worked out in the Ph.D. thesis of Olga Grinchtein [G08].

-- Industrial Liaison --

Working meeting on Integrated Modular Avionics:
On November 12-13, 2007, an ARTIST2 meeting on IMA (Integrated Modular Avionics) was co-organized by Albert Benveniste (INRIA), Paul Caspi (Verimag), in close cooperation with John Rushby (SRI), and hosted by Alberto Ferrari (PARADES) in Rome, Italy. The workshop has gathered participants from aeronautics industry, including manufacturers (Airbus, Boeing, Dassault-Aviation), system suppliers (Honeywell, Wind River), service companies (WW Technology group), labs (SRI, SAE AADL Committee), and academics (TU Vienna, Verimag). Detailed minutes are available from ARTIST2 Web site. More about conclusions are in Section 2.4.5.

Interaction with the Automotive industry
The integrated project SPEEDS has developed a layered meta-model of heterogeneous rich components (HRC) and standardized approaches for the integration of commercial industry standard modeling tools to assemble system-level design models with rich interface specifications by combining models expressed in any authoring tool compliant to the integration standard, including Matlab-Simulink/Stateflow, Rhapsody, and Scade. It is currently integrating a range of analysis methods supporting interface compliance testing and dominance analysis between contracts expressed in an extended automata model.

On March 4, 2008 a *SPEEDS Automotive Day* was organized to discuss with the automotive industry how the AUTOSAR methodology can be supported by SPEEDS technologies, striving to reconcile the advantage of early system-level analysis with the overall AUTOSAR objective of decoupling function design from its implementation. The discussion has been deepened in bilateral meetings between OFFIS and individual automotive companies (May 28: BMW, Sept 3: Bosch and ETAS; planned: Nov 6: Continental). The interaction with the automotive industry will be continued both through direct participation of Artist2 Members in Autosar (OFFIS, CEA), as well as through projects launched through EICOSE.

On June 16, 2008, a SPEEDS tutorial was held in the context of the INCOSE'08 6th biennial European systems engineering conference in Utrecht (http://www.incose.org/symp2008).

Since 2004, CEA has been strongly involved in setting up the Num@tec Automotive working group of System@tic Paris Région competitiveness cluster. In this context, a platform for component based development of automotive system has been launched in September 2007: the EDONA platform (www.edona.fr). It targets tool integration under the AUTOSAR standard and covers both requirements engineering, software architecture design and model based validation.

Interaction through EICOSE
Specifically related to the topic of component based design are strategic initiatives taken by the Artemis Innovation Cluster on Transportation, EICOSE, to create a reference technology platform for embedded systems design, which in particular will strive to harmonize major existing initiatives for component models in embedded systems design, taking into account industry standards such as Autosar and SysML. Eicose has launched project proposals both in the context of ITEA (with an emphasis on open-source developments) as well as the Joint Undertaking Artemis (with an emphasis on the safety critical embedded systems market) towards theses objectives, thus contributing to the overall Artemis objectives of driving future European Standards for Embedded Systems design. This initiative will in particular benefit from the Artist2 activity on component based design for heterogeous systems, as well as on results of related research projects such as SPEEDS and COMBEST.

2.2.2 Adaptive Real-time Cluster

This cluster is composed of the following activities:

A common infrastructure for adaptive Real-time Systems (Platform)
Led by Giorgio Buttazzo (Scuola Sant'Anna - Pisa)

Partner teams (leaders): Giorgio Buttazzo – Scuola Superiore S. Anna (Italy), Luis Almeida – University of Aveiro (Portugal), Gerhard Fohler – Technical University of Kaiserslautern (Germany), Michael Gonzalez Harbour – University of Cantabria (Spain), Alan Burns – University of York (UK), Eduardo Tovar – Polytechnic Institute of Porto (Portugal)

Affiliated teams (leaders): Ivo De Lotto – University of Pavia (Italy), Paolo Gai – Evidence s.r.l. (Italy), Lucia Lo Bello – University of Catania (Italy), Pau Marti – Universitat Politècnica de Catalunya (Italy)

Overview: A research platform for real-time systems is composed of competencies, resources, and tools targeting at the development of control applications with performance and timing requirements.

A shared research platform is essential for experimenting new real-time software technology, including novel scheduling algorithms, resource management techniques, energy-aware policies and overload handling approaches to increase robustness and predictability. Platforms are also used as the basis for transfer research results to industry, as they allow teaching practical knowledge of the concepts and techniques.

Two platforms have been used within ARTIST2 to share competences and test methodologies for real-time operating systems:

1. SHARK (Soft and HArd Real-time Kernel) is an open source modular operating system developed at the ReTiS Lab of the Scuola Superiore Sant'Anna of Pisa running on Intel-based personal computers. It integrates novel real-time algorithms resulting from long-term research efforts. The kernel modularity allows the user to replace scheduling and mutual exclusion mechanisms without changing the application code. This was the main motivation for adopting this kernel. Shark is compliant with the POSIX standard interface.

2. ERIKA Enterprise is a minimal real-time kernel developed by Evidence s.r.l. for small microcontrollers with severe resource constraints. Erika is compliant with the OSEK standard interface, which makes it suitable for developing portable applications, especially in the automotive domain. It runs on a variety of platforms, including Microchip dsPic, ARM 7, Atmel AVR, H8, and Altera NIOS II. ERIKA implements innovative scheduling algorithms such as Fixed Priority with preemption thresholds, Stack Resource Policy (SRP), and Earliest Deadline First (EDF). The kernel is available in double licensing, both GPL and commercial, and has a minimum footprint of 800 bytes of code.

The Shark operating system was mainly used in the first three years of the project to test the effectiveness of novel scheduling and resource management algorithms, whereas the Erika kernel was adopted in the last years for experimenting the development of real-time applications on top of typical embedded platforms with limited resources, like memory, processing power, and energy.

As a shared platform for the Erika kernel we selected the FLEX board (developed by Evidence s.r.l.), which is a Microchip dsPIC evaluation board for embedded applications. The compact design of FLEX makes it suitable not only for development purposes, but also for a direct deployment in the working environment.

RT-Druid is the development environment for ERIKA Enterprise. Based on Eclipse, RT-Druid allows writing, compiling, and analyzing an application in a comfortable environment.

Work in Year 1

Initial definition of the operating system and network features. *The SHARK operating system developed at the Scuola Superiore Sant'Anna of Pisa has been identified (for the reasons explained in* Deliverable 2-2 JPIA-a-ART-Y1*) as the most suited kernel for building a common infrastructure to perform advanced experiments on real-time systems.*

Work in Year 2

Deployment of a working platform for experimenting RTOS and network development. *The SHARK operating system was upgraded according to the partners' needs and deployed on each partner site. A specific workshop has been organized in Pontedera (Pisa) to teach partners how to use the kernel for writing a real-time application and how to write new scheduling and resource modules.*

Work in Year 3

Extensive testing was performed on Shark to identify algorithms and tools to support adaptive RT systems. Specific applications were developed under Shark by the cluster members. Examples are listed below.

- Ball balancing. This control application has been developed by the Scuola Superiore Sant'Anna, in collaboration with the University of Pavia (for sensors and actuators interfacing), Evidence (for kernel support), and Lund (for controller implementation). A two-degree of freedom device has been built using two servomotors and a camera has been used to track the position and velocity of a ball moving on the plate.
- Inverted pendulum. This application has been developed by University of Catalonia (affiliated to Pisa). Shark was used to perform sensory acquisition, actuation, control and to test overload management strategies.
- Mobile Robots. This application has been developed by University of Aveiro. An inverted pendulum has been mounted on a robot car, which has been controlled to keep the pole in a vertical position.
- Overload management techniques. A number of multitask real-time applications have been developed at the Scuola Superiore Sant'Anna to test the behaviour of real-time systems under overload conditions. *Resource Reservations* and *Elastic Scheduling* techniques have been evaluated to cope with transient and permanent overload conditions. *Resource Reservations* techniques basically isolate the temporal behaviour of a task (or subset of tasks) protecting the rest of the systems from potential overruns. On the other hand, *Elastic Scheduling* provides an effective solution to cope with permanent overload conditions. According to this method, task utilizations are treated as flexible springs that can be compressed (by enlarging periods) to reduce the load up to a desired value. Such novel techniques (not yet available in commercial operating systems) have been implemented into Shark as basic scheduling modules to be tested and evaluated in actual control applications.
- Feedback Scheduling in SHaRK**: a First Approach**. Research Report ESAII-RR-06-13, Automatic Control Department, Technical university of Catalonia, July 2006 (http://www.upcnet.es/~pmc16/shark06.pdf). This work reports basic modifications done in the shark kernel to facilitate the application of existing feedback scheduling results. No particular application was implemented. Unofficial modification of the S.Ha.R.K.

kernel available for download (http://www.upcnet.es/~pmc16/all.tar). Unofficial Knoppix ISO with S.Ha.R.K. available for download (http://www.upcnet.es/~pmc16/knoppix.iso). Authors: Josep Guardia, Pau Martí, Manel Velasco and Rosa Castañé.

- Porting of HOLA-QoS. The group at UPM is continuing the porting of HOLA-QoS on top of SHARK. Up till now, the lower levels (Resource Manager) have been ported and a new version of the quality manager (higher layers) is being developed. When it will be ready, the software will run on top of SHARK, so that the negotiation and optimization features of HOLA-QoS can be used in SHARK applications.
- Cibermouse client. University of Aveiro is developing a Shark client for the CiberMouse@RTSS2006 students design competition. The code can be found here: http://www.ieeta.pt/~lau/web_ciberRTSS/tools.htm
- Video processing. TUKL is using Shark as a platform for video processing applications, in particular adaptive resource management for user quality. Algorithms for stream adaptation have been implemented and evaluated on Shark. Representatives from TUKL participated in training for Shark.
- Education at TUKL. TUKL is using the Shark kernel in undergraduate education. Two labs have been: one in which students develop a scheduling algorithm, which they analyse theoretically and practically as implementation on an operating system, i.e., Shark. In the other lab, students implement a simple video processing algorithm which they implement on Shark, learning implementation and overhead issues.
- Traffic Smoothing Techniques. University of Catania (affiliated to Pisa) is implementing Traffic Smoothing Techniques on the Shark OS, to make experiments on distributed real-time systems.

At the same time, a new platform was developed in the third year for experimenting the development of real-time applications on top of typical embedded platforms with limited resources. The new platform, consisting of the FLEX board and the Erika Enterprise kernel, is described in Section 1.4.

Final Results

In Year 4, we concentrated on the development of an embedded platform for running real-time applications under severe resource constraints. The following control applications have been developed using Erika Enterprise as a real-time kernel and Flex as a hardware platform.

- Ball and plate balancing. A two-degrees-of-freedom balancing device has been built at the Scuola Superiore Sant'Anna to control the trajectory of a ball on a plate actuated by two servomotors. The position of the ball is detected by a resistive touch screen mounted on the plate. http://www.evidence.eu.com/content/view/276/266/
- Visual tracking. A visual tracking application was developed at the Scuola Superiore Sant'Anna, where a moving ball was followed by a mobile CMOS camera using the FLEX board. The CMOS Camera is capable of returning JPEG images to the connected FLEX board hosting a Microchip dsPIC. The Flex Board also controls two servomotors which are used to articulate the camera, thereby maintaining the focus on the rolling ball. http://www.evidence.eu.com/content/view/277/266/
- Inverted Pendulum at SUSPI. An inverted pendulum was controlled using the FLEX board with Scilab/Scicos at SUPSI (Scuola Universitaria Professionale della Svizzera Italiana), Lugano, Switzerland. The FLEX Base Board and the FLEX Multibus Board with a CAN module were used for swinging-up and maintaining the inverted equilibrium. The Source code was entirely generated using Scilab/Scicos, an automatic code

generator for control systems.
http://www.evidence.eu.com/content/view/274/266/

- Inverted Pendulum at UPC. An inverted pendulum was implemented using the FLEX board and Erika by the Distributed Control Systems group at the Automatic Control Department, Technical University of Catalonia, Barcelona, Spain. The FLEX Board and the FLEX Multibus Board with a RS232 module was used for both swinging-up and maintaining the inverted equilibrium.
http://paginespersonals.upcnet.es/~pmc16/08EvidenceNoteRTpend.zip
- DC Motor control. A DC servomotor was controlled using the FLEX board with Scilab/Scicos at SUPSI (Scuola Universitaria Professionale della Svizzera Italiana), Lugano, Switzerland. The FLEX Base Board and the FLEX Multibus Board with a CAN module were used for Servo control of a DC Motor. The Source code was entirely generated using Scilab/Scicos, an automatic code generator for control systems.
http://www.evidence.eu.com/content/view/273/266/
- Hexapode robot control. An 18-DOF hexapod robot was completly designed and developed at the University of Florence by Andrea Foschi in 2005. It was later tamed by Marco Natalini and Alessandro Mambelli using Evidence Srl's FLEX Light board and ERIKA kernel. The main purpose for adopting FLEX was due to its low-cost development kit that permits easy addition of features, i.e., sensors and behaviour. Since then, a number of students have worked on this hexapod.
http://www.evidence.eu.com/content/view/261/266/
- Educational experiments. An embedded control system was developed with the Erika+Flex platform at the Automatic Control Department of the Technical University of Catalonia (Spain), with the purpose of setting a laboratory experiment for educational purposes. A real-time control of dynamical system was designed to drive students to a better understanding and integration of the diverse theoretical concepts that often come from different disciplines such as real-time and control systems.

QoS aware Components (NoE Integration)
Led by Alejandro Alonso (Universidad Politécnica de Madrid)

Partner teams (leaders): Alejandro Alonso (UPM), Jean-Marc Jezequel (INRIA), François Terrier (CEA), Jacques Pulou (FTR&D).

Affiliated teams (leaders): Laurent Pautet (ENST), Stefan van Baelen (K.U. Leuven), Marisol García-Valls (U. Carlos III of Madrid), Virginie Watine (Thales).

Overview: QoS concepts are starting to appear in component standards, but are far from mature. Partners in this JPRA have expertise in different aspects necessary for progress. An example is the request for proposal at the Object Management Group that is currently demanding solutions for the integration of some QoS facilities in CORBA Component Model.

There are a number of techniques and methods required for the industrial use of QoS aware components such as:

- Notations for the description of components models including functional and QoS (also know as non-functional) aspects. The integration of this information in the interfaces is of primary importance.
- Automatic generation of analysable models from the UML model.

- Composition mechanisms for determining whether the interconnection of two components is feasible and for deriving the non-functional characteristics of a group of connected components. This work is related with the adaptation of component execution to changes in the environment.
- Component frameworks to support the runtime composition of QoS aware components.

Work in Year 1

Some partners cooperated in the development of the OMG standard "UML profile for QoS and Fault Tolerance", which was finally approved on May 2005.

The main result of this work was the concrete identification of the more concrete integration topics and the start of this work. This final job was done during a meeting that allowed the partners to know each other and discuss their interests. The identified integration topics were:

- Consistent alignment between the QoS modelling style of MARTE (with basis on Schedulability, Performance and Time) and that of the UML Profile for QoS and Fault Tolerance. The fist one is mainly related to time and performance aspects, while the second is more general, as it tries to provide means for specifying any other QoS characteristic. Partners involved: CEA, Inria, UPM.
- With respect to composability, the interest is focused upon the development of a contract model with well-founded semantics with respect to time and execution. This contract model handles (some) QoS characteristics. Partners involved: CEA, Inria, UC3M, UPM
- Finally, the support for the execution of QoS aware components requires components infrastructures with this support. UPM (QoS in the Robocop framework), UC3M, CEA and Thales (CCM based extensions) have done previous work on this topic. They have also proposed containers to simplify components development. The goal will be to interchange the approaches to try to get their particular merits and to propose new concepts for their future evolution.

The work on these topics has started during the previous work period.

Work in Year 2

The first issue was the identification of notations for the integrated description of functional and QoS properties in general component models. Some partners participating in this activity are active on two OMG standardization efforts that define such notations: OMG standards on UML profiles on "Real-Time and Embedded systems modelling and analysis (MARTE)" and "QoS and Fault Tolerance". A complementary activity is the definition of catalogs of QoS attributes of a QoS characteristic, in order to try to develop techniques for their modelling. In the "QoS and Fault Tolerance" UML profile it is defined a general catalogue. In addition, QoS attributes for safety have been defined. An activity started during this period of time was the selection of a case study and QoS attributes, model them with different profiles or techniques and compare which is the most suitable in each case,

An important advantage of modelling QoS properties is the possibility of generating automatically models that can be used as input for analysis tools. The suitability of the previously mentioned profiles and attributes has been a subject of work for this year, with special focus on the QoS properties for timing and safety.

Another issue was to define the composition of QoS aware components. In this case, the connection of two components is only feasible if the provider includes the required operations with the proper functionality and QoS features. The common approach is the definition of a

contract model where the specificities of the functions to be provided are determined and that serves as the basis for the evaluation of the feasibility of the composition and the resulted quality.

There are components infrastructures that provide the required support for the execution of components. However, there are no mature infrastructures supporting QoS aware components. In addition to the general functions, support for the negotiation between components, for finding a suitable provider, and with the system, for getting the resources required for the system execution are needed. There are some initial works towards these goals; such as extensions to CCM by Thales and CEA and the QoS support at Robocop done by UPM.

Work in Year 3

The first issue was the identification of notations for the integrated description of functional and QoS properties in general component models, in order to reason about whether a component or a set of them fulfils a certain set of requirements. The integration of non-functional aspects allows for ensuring this property along the development lifecycle, hopefully, in an automatic way. Partners in this activity have continued their efforts in the standardization of this type of notations in the OMG (*Object Management Group*). The work on the "*UML Profile for QoS and Fault Tolerance*" standard has been subject of a revision considering a number of issues that were submitted by users. As a result, a new version of the standard has been approved on December 2006. The efforts on the "*UML Profile for Modelling and Analysis of Real Time and Embedded Systems*" (MARTE) have concluded with its adoption on June 2007 by the OMG. In addition, these profiles have been used to model safety and time.

The automatic generation of models for analyzing a particular characteristic is one of the major advantages of modelling together the system functional behaviour and QoS. Early evaluations of the system could guide to a better and cheaper end-product. It is intended to do them over the architectural designs; however, these designs are evolving until the architecture is completely specified. Automating such analyses in this changing environment is of great importance to aid engineers. Time and safety are the QoS characteristics that have been modelled in the context of this work. As a natural extension, the automatic generation of models for analysing them have been tackled.

The definition of the composition of QoS aware components has also been subject of research during this year. A UML profile for this job is under development. This work has naturally added a research topic: adaptability in QoS-aware systems. When composing a set of components, it is necessary not only to provide the required functions but the QoS characteristics as well, for setting the contracts. It is also required to determine which is the quality provided by this set as a whole. In addition, if these components can offer or require functions with different quality levels, it is of interest to know for each possible combination its feasibility and the overall provided quality. This information is relevant to statically evaluate the adaptability of a system and to change on runtime the provided quality according to the execution context.

The work on QoS management facilities in component infrastructures has been refined during this period of time. Additional functions have been added and a number of programming errors have been fixed. In addition, the API of the Robocop QoS Manager has been the basis for the specification of the resource and quality management in the ISO/IEC 23004 standard, which is specially suited for embedded systems. The integration of QoS in CCM also attempts to provide this type of runtime support.

Final Results

These achievements are aligned with the four main research lines:

a. Specification of QoS properties using UML profiles and aspect-based approaches
b. Generation of analysable models from the UML models
c. Composition of QoS-aware components and adaptability
d. QoS support in run-time components frameworks

For each technical achievement, there is an indication of the activity to which is mainly related.

-- Characterisation of services --

UC3M has developed a characterisation of QoS properties for service-based applications to enable functional composition in networked embedded systems. Contributions to the composition of applications in real-time have been made by developing algorithms that are able to make calculations with respect to selected parameters. This work has been done in collaboration with the University of Aveiro and UPM.

The solution for composing services-based applications developed by Uc3M is based on the definition of figures of merit, each one corresponding to a set of composition criteria, and the development of composition algorithms. The composition algorithms selects schedulable sets of implementations, while the figures of merit discriminates between different paths, in order to select the best according to a QoS criteria, e.g. minimization of utilization factor, minimization of the response time of the whole application, etc. On the other hand, UC3M developed two composition algorithms: an exhaustive one, suitable for off-line composition and an improved one, based on heuristics, suitable for on-line composition.

-- QoS in Component-Based Approaches --

Important results [1, 2, 3, 8] were obtained regarding *fault-tolerance* and *adaptability* in the INFLEXION project (Adaptable and Flexible Execution Infrastructure) from the “Usine Logicielle” program (SYSTEM@TIC PARIS-REGION Cluster). Replication can be achieved in a transparent way within a component-based approach. Replication is declared statically on model level and supported by middleware layer. It makes use of connector extension eC3M, www.eC3M.net.

In the context of Flex-eWare (French ANR project), CEA LIST defines a common meta-model for component aspects embedded and real-time applications, particularly oriented to UML-based models. Initially, CEA LIST carried out an empirical study that provides an intuitive description of composite structure semantics. Among practical solutions, a mechanism of encapsulating explicit behaviours in component ports has been proposed [4].

-- Consolidation of the MARTE standard --

One key standardization effort is the UML™ profile for Modeling and Analysis of Real Time and Embedded Systems (MARTE) at OMG. Currently, MARTE (http://www.omgmarte.org/Documents/Specifications/08-06-09.pdf) is in a second phase of the finalization tasks (version Beta 2) and is planned to be released in March 2009 (version 1.0). Regarding QoS and component-based support, some modelling features are being aligned to SysML and other related standards (AADL, Autosar, EAST-ADL, among others) [5]. At the same time, a more precise semantics for VSL (Values Specification Language) is being provided. VSL is the expression language that is the basis to specify non-ambiguous non-functional aspects in MARTE.

CEA LIST is participating in a number of projects that are providing tool support for MARTE. Among them, FP7 INTERESTED, ANR Lambda and Usine Logicielle are specially interested in providing QoS modelling support in embedded systems development **Erreur ! Source du renvoi introuvable. Erreur ! Source du renvoi introuvable.** [6, 7]. Thales is collaborating with Thales in the context of this project.

Thales and Universidad de Cantabria have also participated in the standardization process of the MARTE standard.

-- Generation of Analyzable Models from the UML Models for Safety --

Safety-critical software requires integrating verification techniques in software development methods. Software architectures must guarantee that developed systems will meet safety requirements and safety analyses are frequently used in the assessment. Safety engineers and software architects must reach a common understanding on an optimal architecture from both perspectives. Currently both groups of engineers apply different modelling techniques and languages: safety analysis models and software modelling languages.

UPM has developed solutions to integrate both domains coupling the mentioned types of notations. A model-driven development approach and the use of a platform independent language are used to bridge the gap between safety analysis (failure mode effects and criticality analysis and fault tree analysis) and software development languages (e.g. unified modelling language). Language abstract syntaxes (meta-models), profiles language mappings (model mappings) and language refinements, support the direct application of safety analysis to software architectures for the verification of safety requirements. Model consistency and the possibility of automation are among the main benefits. Safety annotations are included in the traditional software models and tools have been developed to automatically extract safety models that are a direct input to safety analysis tools. During the last year the previous work has been refined and extended with more rich safety modelling capabilities.

-- Task Force of "MDA Tool Component RFP" --

In June 2006 the OMG raised a call for proposals for the standard: "MDA Tool Component RFP" (MDATC). Its goal is to create modelling tool support for defining a packing mechanism and interchange of development artefacts based on MDA, and to allow the reuse of development support tools based on MDA. This call can be found in: http://www.omg.org/docs/ad/06-06-09.pdf

UPM, along with a number of other companies and universities (Softeam, Thales, Universidad de York, France Telecom, and Adaptive) developed an initial proposal to meet the requirements of this standard call, which was presented on June 2007. (http://www.omg.org/docs/ad/07-06-12.pdf, http://www.omg.org/docs/ad/07-06-04.doc). UPM participated with more than a 50% of the original proposal, in particular, in the definition of the MDATC meta-model and a RAS (Reusable Assets Specification) profile.

OMG members studied this proposal and defined how to update the proposal. UPM is involved in the development of a reviewed proposal that was presented on September 2008 (http://www.omg.org/docs/ad/08-09-07.pdf). The group is currently developing a final version, taking as the basis the already presented proposal and comments from different reviewers. UPM is experimenting with this proposal for creating a package with the developed safety modelling tools, in order to facilitate the distribution of these tools.

-- Composition of Quality-Adaptable Components--

Quality of service adaptability refers to the ability of components/services to adapt in run-time the quality exhibited. A composition study from a quality point of view would investigate how

these adaptable elements could be combined to meet some system quality requirements and thus assess its feasibility. Enclosing quality properties with architectural models has been typically used to improve system understanding. Nevertheless these properties along with some supplementary information about quality adaptation would allow us to carry out a composition study during the design phase and even to predict some features of the adaptability behaviour of the system. Analogous supplementary information could be used to denote that several alternatives (e.g.: vendors, implementations) concerning quality can be used even when the system is not run-time adaptable, and a similar composition study would try to select the best choice. Existing modelling languages and tools lack enough mechanisms to cope with adaptability, e.g. to describe system elements that may offer/require several quality levels.

UPM has developed an approach that allows the reuse of existing modelling languages and tools, combine them and create new ones to tackle the problem of quality of service adaptability and composition: extending the standard UML profile for QoS to define adaptable elements, using the OCL language to define QoS functions, applying MDD transformations to process QoS metadata, adapting OCL tools to evaluate QoS expressions, solving the composition problem with search algorithms. The final goal of this work is to evaluate architectural models to predict system′s QoS behaviour before it is implemented.

-- QoS support in run-time components frameworks--

MPEG, a working group in ISO/IEC, is currently working in the standardization of an Application Programming Interface (API) for Multimedia Middleware (M3W), that as explained in the introduction, will allow application software to execute multimedia functions with a minimum knowledge of the inner workings of the multimedia middleware as well as to support a structured way of updating, upgrading and / or extending the multimedia middleware.

The Application Programming Interface mentioned above and a realization technology is specified in detail in ISO/IEC 23004 part 1-7. ISO/IEC 23004 part 8 is the reference software provided for this standard. Part 4 defines and API for resource and quality management. UPM has participated in this part since its beginning. The API provides the basic means to allow QoS-aware components to notify its quality information, to search for components providing a given quality and compose quality information. It is based on the HOLA-QoS quality management middleware.

This standard will include a reference implementation of this API to let final users experiment with the proposed API. UPM developed an initial version of a reference implementation for part 4 that includes a quality and a resource manager.

During the last reporting period, UPM has upgrading the previous version of the quality and resource manager initial version to make it compatible with a recent new version of a reference implementation of the components runtime environment and associated tools. In addition, it has been developed a set of components that encapsulate these managers and that allows for an easier modelling of the users QoS-aware components within this framework.

-- Rule-based approach to model adaptation policies --

In spite of new methods and technologies in software engineering such as CBSE, or AOP, it is still difficult to talk about adaptation since adaptation policies might impact the architecture, the configuration data, and some extra-functional features as well.

During 2008 the Inria team has finished the design of its QoS management platform. This platform uses a model driven engineering process where the designer can define adaptation policies for QoS properties of the system under design. The process and the tools rely on the metamodelisation of QoS properties, monitoring properties, adaptation policies, and

simulations. The tools can also simulate the system to study the impact of changes in policies. A tool then generates monitors that supervise the running system to detect policy violations and activate adaptation behaviours. The work has been presented at the Embedded Real Time congress **Erreur ! Source du renvoi introuvable.** at the French conference on Object Oriented Design LMO **Erreur ! Source du renvoi introuvable.**.

Flexible Resource Management (Cluster Integration)
Led by Gerhard Fohler (Kaiserslautern)

Partner teams (leaders): Gerhard Fohler (Technische Universität Kaiserslautern), Giorgio Buttazzo (Scuola Superiore S. Anna), Michael Gonzalez Harbour (University of Cantabria), Luis Almeida (University of Aveiro), Eduardo Tovar (Polytechnic Institute of Porto), Alejandro Alonso (UP Madrid), Alan Burns (University of York (UK).

Affiliated teams (leaders): Marisol García-Valls - U. Carlos III, Madrid), Liesbeth Steffens (NXP), Ivo De Lotto – University of Pavia (Italy), Lucia Lo Bello – Univ. of Catania (Italy).

Overview: In some application domains, such as multimedia, applications are very expensive in terms of resource consumption. In other applications domains, such as automotive, mobile telephony or even building automation, the resources are scarce and there is a growing pressure to integrate resources even further and optimize their use. In both cases, timeliness directly relates to user perceived quality, e.g., smoothness of the video stream. Furthermore, efficient resource usage is key issue not only for cost considerations, but also for competition on a feature bases: better resource usage – more features.

Both resource demands, e.g., MPEG-2 video streams, and resource availability, e.g., available bandwidth on wireless links, fluctuate rapidly and unpredictably; worst case assumptions will lead to extreme over provisioning. Consequently, methods for adaptive resource management are required.

Trading resource usage (processing, communication and memory/storage, inter-device and intra-device) against offered output is known as QoS (Quality of Service). The different resources cannot be considered separately, interferences and inter-resource tradeoffs have to be taken into account because they affect the application output. The tradeoffs have to be made at different time scales, in order to match the time scales of the system dynamics.

Theory for independent scheduling algorithms is well defined in the areas of event triggered and time triggered systems, but few theoretical results have been achieved in trying to integrate these approaches. Some partial results exist for simplified architectures, but it is necessary to enhance them by taking into account all of the requirements of modern real-time systems including distributed ones. In addition to the development of theory, a framework needs to be built in order to allow a flexible way to handle different scheduling algorithms for different kinds of resources, and evaluate their applicability to real application domains.

Work in Year 1

In the first period, the technical results were achieved in the following areas: video stream demand analysis, identification of scheduling algorithms and kernel mechanisms for stream adaptations based on integrated, flexible scheduling; adaptive resource management for network bandwidth management, multi resource management, in particular with respect to cache aware scheduling; middleware support for QoS management.

Furthermore, the ART cluster has been in active contacts with relevant industry to gather understanding of realistic requirements and to identify research topics and baselines relevant

for industrial and academic research. Partners has been giving presentations at the Philips Software Conference – Real-time Workshop and had meetings with Nokia, Ericsson mobile platforms and Visual Tools from Spain. The goal has been to go as far as possible towards the actual engineers for better understanding and prepare for a specific industry – academia workshop with selected participants.

New scheduling mechanisms for integrating overload management techniques with energy-aware strategies were investigated in the context of real-time systems. The new scheduling mechanisms were analysed to guarantee timing constraints while minimizing energy consumption, and a kernel infrastructure was developed into the Shark operating system in order to facilitate their implementation. Moreover, the assessment of the (m,k)-firm model was done for its implementation over the FTT-Ethernet protocol.

Work in Year 2

-- Temporal Constraints for Video streaming --

Philips and TUKL have studied temporal constraints of video streaming. As sources for the constraints we looked into semantic stream dependencies from MPEG decoding, as well as the temporal impact of devices and their resources in the end-to-end delivery chain of a stream. The work was carried out with industrial partners in the area. The results have been fed into other activities in the cluster, in particular w.r.t. to scheduling and networking.
http://rts.eit.uni-kl.de/research/mediaprocessing

-- Integrated real-time scheduling and cache management --

Philips and TUKL continued work on integrating real-time scheduling and cache management on multiprocessor platforms. To this end, we carried out experiments to study cache behaviour on the actual platform and formulated a number of scenarios with increasing complexity. A joint PhD student has carried out the work.

-- Adaptive service configuration for Quality-of-Service aware middleware --

In order to support dynamic services with adaptive QoS requirements, we proposed a dynamic scheduler which is able to react to load variations. Isolation between different services is still achieved by guaranteeing a minimal service quality to accepted services and by an efficient overload control that considers the challenges and opportunities of dynamic distributed embedded systems. This scheduler was also extended taking into consideration simple dependencies between services' QoS attributes.

-- *Server Based Flexible Scheduling --*

Schedulability analysis techniques were developed for server-based systems that can be used to schedule different kinds of flexible timing requirements, such as those needed to integrate control systems with multimedia activities. In particular, this work focused on hierarchical scheduling analysis and design techniques A further issue was the dimensioning of the parameters of a server for minimizing the average response time of the served activities. A statistical approach was addressed in order to compute the probability of missing a given deadline. Partners were SSSA, Cantabria, TUKL of the cluster and the partners of the FIRST and FRESCOR EU STREP consortia. www.frescor.org

-- Adaptive resource management for networks --

Work concerned the analysis of the achievable QoS guarantees in wireless networks. The achievable end-to-end QoS guarantees were investigated as a function of the guarantees provided by the underlying resource schedulers. Further activities dealt with network protocols to efficiently support dynamic bandwidth management with strict QoS guarantees in Ethernet-based networks, awireless time-token communication protocol that allows providing real- time guarantees for real-time messages and tune the allocated bandwidth according to the required QoS was developed. Aveiro, Porto, SSSA, and TUKL carried out work. http://www.hurray.isep.ipp.pt/activities/art-wise/
http://rts.eit.uni-kl.de/research/mediaprocessing

-- Adaptive service configuration for Quality-of-Service aware middleware --

An iterative refinement approach with the ability to trade off deliberation time for the quality of the solution was specified. The work also addressed the problem of dynamically changing system conditions, allowing the system to make QoS adaptation decisions in response to fluctuations in the nodes service load, under the control of the user. Monitoring the stability period and resource load variation of Service Level Agreements for different types of services was used to dynamically adapt future stability periods, according to a feedback control scheme. Work was done by Madrid and Porto (www.hurray.isep.ipp.pt/activities/qos).

Cluster partners have developed kernels that provide these facilities and, hence, could be suitable to act as the lower layer of a HOLA-QoS based system. Work that it was under development was to port HOLA-QoS on top MARTE (Cantabria) and SHARK (Pisa) kernels. One result of this work is the possibility of experimenting with the adaptation techniques that these advanced resource kernels provide. Some publications on HOLA-QoS can be found at http://www.dit.upm.es/str .

-- Resource availability prediction --

The resources typically used in-home entertainment applications (e.g., video/audio streaming) exhibit fluctuating availability. It is desirable to have mechanisms for indicating the available bandwidth during system runtime.

A comparative analysis of bandwidth estimation techniques for WiFi links has been carried out. In particular, the analyzed estimation techniques include several statistical and control-based algorithms. The analysis has identified the best suitable techniques taking into account the specific behavior of WiFi links. Work was carried out by UPC and TUKL. Analysis available at http://www.upcnet.es/~pmc16/nde_06.pdf .

Requirements for integrated-resource scheduling framework
A workshop on “Requirements for Flexible Scheduling in Complex Embedded Systems” was held in Massy (Paris) in June 2006, with the objective of developing a set of requirements for building a flexible scheduling framework for applications demanding various types of tasks, constraints, and scheduling paradigms within the same system, and paying attention to the integration of multiple resources. The workshop was very sucessful and brought together 20 participants.

Baselines for integrated-resource scheduling framework
The FIRST (Flexible Integrated Real-Time Scheduling Technologies) IST project that finished in 2005 produced as its main result a contract-based scheduling framework, called FSF that was capable of scheduling multiple application components with various kinds of requirements for CPUs and, to a limited extent, for networks in distributed systems. This framework was

selected as the baseline for the more ambitious framework that is being developed in this activity and that will take into account the integrated scheduling of multiple resources.

The FTT framework, in which a master manages the synchronous activities in a distributed system or cluster, was also extended to micro-segmented switched Ethernet-based distributed systems, having revealed potential to provide efficient support to the contract model, to dynamic QoS management and to integrated resource scheduling in distributed environments. The VTPE protocol (Virtual Token-Passing Ethernet), which supports event-triggered communication with real-time guarantees and high bandwidth utilization, was also extended with appropriate mechanisms to support isochronous traffic, more adequate for some applications, e.g. multimedia transmission.

New theoretical developments
The contract-based scheduling framework needs to be implemented using a specific scheduling strategy, and the most effective approach for this case is the server-based hierarchical scheduling in which an application or application component is scheduled over a protected bandwidth-preserving server (such as a periodic server, a sporadic server, or a constant bandwidth server) and individual threads in that component are scheduled by a higher-level scheduler that uses the bandwidth provided by the server. Theory was developed towards being able to analyze such scheduling schemes.Work was done by the University of York together with Technische Universiteit Eindhoven (TU/e) on the analysis underpinning the use of CAN in real-time systems.

SSSA developed the following theoretical results: energy-aware scheduling algorithms for processors with dynamic voltage scaling and discrete frequency levels; a method for minimizing the deadline of periodic tasks with the objective of reducing delay and jitter; a general methodology for performing sensitivity analysis of fixed priority periodic systems with configurable periods and computation times, allowing the system designer to derive the feasibility region of a task set and compute the maximum parameter variations that keep the system feasible

Work carried out at the University of Aveiro also exposed a couple of anomalies related to the definition of critical instant in hierarchical scheduling scopes found in communication systems that led to optimistic worst-case response time analysis in the past. Adequate methods were devised to cope with such anomalies.

The Polytechnic Institute of Porto provided new theoretical developments on: a new multiprocessor scheduling approach with a higher utilization bound and with few preemptions, able to trade the utilization bound for preemptions; new flexible admission control algorithms for IEEE 802.15.4 networks improving the bandwidth utilization compared to the explicit allocation used in the IEEE 802.15.4 protocol; a new server-based scheduling approach for handling isolation and overload control on distributed cooperative systems;

Flexible architectures and communication protocols for networks used in distributed embedded real-time systems
This research was partly done in collaboration between the following ARTIST 2 partners: University of Pavia, University of Catania (affiliated partner) and Malardalen University, Sweden. It consists of the following activities:

- Integration of networked subsystems in a resource constrained environment.
- Facilitating subsystem integration by decoupling priority and identifier in CAN messages.
- Interconnection of real-time networks in factory automation and in the automotive domain.

- Design issues and transmission protocols for wireless networks used in factory communication
- Bluetooth to support real-time traffic in factory communication.
- Modelling of wireless real-time communications for land monitoring systems.

In addition, a joint work was done by Pavia, Pisa and Aveiro to support connectivity tracking in mobile ad-hoc wireless networks subject to real-time constraints.

Work in Year 3

The collection of application requirements has been extended beyond the video streaming domain. Activities were carried out in particular together with the FRESCOR project, including a joint requirements workshop and a meeting with the industrial advisory board of FRESCOR on the topic.

Algorithms for the integration of CPU scheduling and cache management have been developed and analysed for their effectiveness on actual boards.

An architecture for flexible functional composition on networked real-time applications was developed.

An architectural model of a flexible scheduling framework has been developed. The framework is capable of handling multiple concurrent activities with different criticality and timing in the same system, integrating the management of different kinds of resources such as processors, networks, memory, energy, and shared objects with time protection.

The flexible scheduling framework has been implemented for different execution platforms and networks.

A number of new theoretical results have been developed.

The feedback control approach has been applied to server-based real-time systems, in order to automatically control the partial utilization of each server and for online adaptation for networks.

The HOLA-QoS and the AquoSA framework have been continued and adapted for use with other kernels by cluster partners.

An architecture to support dynamic service composition has been devised

Methods for soft real-time systems based on stochastic response time analysis has been developed, including for overrun handling.

Final Results

-- Application requirements (all) --

The delayed workshop in the domain of media processing with engineers has finally taken place in the form of a two-day workshop organized by Liesbeth Steffens of NXP at the NXP premises in November 2007. It brought together members of the ARTIST cluster and various NXP/Philips groups from different locations.

-- Integrated CPU scheduling and cache management (TUKL, NXP (formerly Philips)) --

NXP and TUKL have continued work on integrating real-time scheduling and cache management on multiprocessor platforms. We have carried out thorough experiments to study the impact of cache usage and scheduling on predictability on the actual platform. We developed a first algorithm with exhibits encouraging behaviour. Previous work has been extended to include Bus dependencies.

A joint PhD student is carrying out the work. Publications have been submitted. [TUKL1-3]

- Architectural model of the flexible scheduling framework (All) --

An architectural model of a flexible scheduling framework had been developed in previous years and has been extended in the current reporting period. The framework is capable of handling multiple concurrent activities with different criticality and timing in the same system, integrating the management of different kinds of resources such as processors, networks, memory, energy, and shared objects with time protection. The framework is independent of the underlying implementation, and can run on different underlying scheduling strategies. It is based on establishing service contracts that represent the complex and flexible requirements of the applications, and which are managed by the underlying system to provide the required level of service

Based on the experience gained during previous years the contract model has been updated and its API has been redesigned. The main changes made to the contract model have been:

- Updated energy management API to add operations to get more information from the system.
- New disk-bandwidth management services. A very simple module has been added to the framework to explore the possibility of handling contracts related to disk bandwidth.
- New feedback control API. The API of the feedback control module has been totally rewritten, based on the implementation and usage experience.
- Memory management. New services have been defined to be able to distribute spare capacity relative to memory resources..

-- Implementation of the flexible scheduling framework (UC, UPVLC, CTU, SSSA, Aveiro, York, Thales Communications France) --

The main implementation effort in the flexible scheduling framework has been on the Integration of new resources. New communication networks have been added: CAN bus, as a representative of field busses, Wifi, as a representative of wireless networks, wired Ethernet [UC1], and switched Ethernet using industrial switches. An implementation to manage disk bandwidth is underway. Implementations on FPGAs and multiprocessor systems are also underway. In addition, the effects of power management are being introduced in one of the implementations.

The framework has been designed to be implementable on different platforms and work has been carried out to implement the framework on POSIX RTOSs (Partikle/RT-Linux, MaRTE OS), a commercial RTOS used in telecommunications (OSE), and in main-stream Linux kernels. An analysis module developed in the University of York was integrated in the framework to provide advanced schedulability analysis capabilities. [York4]

Work has been performed in cooperation with Thales Communications France (TCF) to Integration the contract-based scheduling framework with a component-based framework. The component based technology is the microCCM framework, that implements the component-container model with an infrastructure that is independent of CORBA. The integration with the

contracts is provided by adding two kind of services: one managing task creation (called ThreadActivationService), and another one managing scheduling attributes (called SchedulingAttributeService). The contracts are declared in the deployment and configuration plan, and therefore the corresponding tool has been modified to read this information and automatically generate the code to create the contracts and manage the interceptors that bind threads to the corresponding contracts. [UC2][UC3][UC4].

The work has been carried out by the members of the FRESCOR EU project, together with other members of the ART cluster, in particular UPM and Aveiro.

-- New theoretical developments (York, SSSA, UC) --

A number of theoretical developments have continued to be made during this last year.

- For hierarchical scheduling, the resource sharing model has extended to include EDF scheduling. [York11]
- For standard fixed priority scheduling, the Multi-Frame extension (to the standard model) has been investigated and exact analysis developed for a number of variations in the model including arbitrary deadlines and release jitter. [York1],[York5]
- For systems that allow online/dynamic admissions, the efficient scheduling algorithms for fixed priority scheduling that were developed and verified last year have now been implemented and tested. These algorithms allow spare capacity to be allocated to applications in an effective and efficient way, thereby enhancing the adaptability of such applications. [York4],[York12]
- Means by which an application can obtained the correct parameters for its virtual resource has been developed. [York2],[York3],[York6],[York9],[York10]
- The application of fixed priority scheduling to wormhole NoE protocols has been investigated. This work includes priority assignment and buffer size reductions. [York7]
- For hierarchical systems that are implemented on an EDF scheduler, necessary and sufficient analysis has been developed that mirrors the fixed priority approach, and thereby allow mixed EDF and fixed priority systems to be deployed. These algorithms, that were initially developed during the previous year, have been improved and show to be optimal (in most cases). [York11]
- As a result of using the POSIX specification of the sporadic server scheduling algorithm in the context of a resource reservation framework, it has been discovered that this specification has a bug that causes that the timing behaviour of the policy is more intrusive than expected on lower priority tasks. Consequently, a new definition of the algorithm that is appropriate for operating systems and real-time networks has been defined and its timing properties have been proven. The new scheduling protocol has been implemented on an implementation of resource reservations on the CAN bus. [UC5]
- Dynamic resource reservations require the use of mode change protocols when new activities arrive at the system, or when changes to an existing resource allocation are needed. An effective mode change protocol that is usable in these dynamic systems is the idle-instant protocol, in which operation in the new mode is deferred until an appropriate instant in the schedule. While this protocol is usable in processing resources, it is very difficult to use in the networks due to the synchronization that would be needed to determine the idle instant. As a consequence a new mode change protocol has been developed that is usable in networks, and performs in a time interval that is not longer than the interval used in the idle-instant protocol. [UC6]

-- Featuring switched Ethernet with flexible traffic scheduling (Aveiro, Valencia, Malardalen, University of Pensylvania, CMU) --

Switched Ethernet is, nowadays, used from office to industrial automation and even in embedded systems. However, COTS switches still present limited traffic scheduling capabilities, normally restricted to FIFO queues and a few fixed priority levels. This limitation makes distributed systems design more complex and less resource efficient or with less guarantees. This problem has been tackled by Aveiro, Valencia and Malardalen by putting together the FTT-SE protocol, channel bandwidth management techniques and server-based scheduling.

FTT-SE allows removing the traffic scheduling of switched Ethernet systems out of the switches and into a specific node. This grants a high flexibility to the traffic scheduling and opens the way to carry out advanced scheduling features such as dynamic QoS management and server-based scheduling. Concerning the former, Aveiro and Valencia (Alcoy group) have been cooperating in the scope of an industrial video surveillance application to setup a dynamic QoS management system that maximizes the channel bandwidth use and the QoS provided to the cameras [AV1]. On the other hand, the FTT-SE protocol was also used to implement a server-based traffic handling approach that increases system robustness with respect to nodes temporal misbehaviours. A preliminary implementation has been reported in [R8].

Dynamic bandwidth management grants a high level of resource utilization while server-based scheduling brings in a high level of isolation thus favouring composability. Moreover, to enforce continued timeliness even during reconfigurations or adaptations, prompt schedulability tests must be used. This issue was also tackled by the Aveiro group, also in collaboration with two US institutions, namely University of Pennsylvania and CMU. Concerning the admission control, or QoS management, adequate utilization-based schedulability tests were developed, namely the Local Utilization Bound [R9] that improves resource utilization, and the adaptation of existing bounds to a model that includes release jitter being thus applicable to the output ports of switches where this phenomenon frequently occurs.

-- Dynamic QoS Adaptation (Porto) --

Due to the growing complexity and dynamism of many embedded application domains (including consumer electronics, robotics, automotive and telecommunications), it is increasingly difficult to react to load variations and adapt the system's performance in a controlled fashion within a useful and bounded time. This is particularly noticeable when intending to benefit from the full potential of an open distributed cooperating environment, where service characteristics are not known beforehand and tasks may exhibit unrestricted QoS inter-dependencies.

Therefore, we have extended the adaptive resource framework, taking into account services' inter-dependencies and quality constraints, with an anytime QoS control policy in which the online search for the best set of QoS levels is combined with each user's personal preferences on their services' adaptation behaviour. We have also taken into consideration shared resources and precedence constraints in the scheduling of the real-time tasks. The concept of bandwidth inheritance is combined with a greedy capacity sharing and stealing policy to efficiently exchange bandwidth among tasks, minimising the degree of deviation from the ideal system's behaviour caused by inter-application blocking.

-- Assessments on architectures and communication protocols for wireless networks used in factory automation (U. of Catania, SSSA) --

This research was done by the University of Catania (affiliated partner) and addressed two main application scenarios, i.e., factory automation and wireless sensor networks. Part of the

activity was done in collaboration with the Scuola Superiore Sant'Anna di Pisa. Joint publications dealing with flexible scheduling mechanisms for network based on well-known wireless standard protocols are in preparation.

The work in [CT1] proposes an approach to overcome some limitations on the way the IEEE 802.11e handles real-time industrial traffic through a dynamic adaptation of the back-off parameters for the different Access Categories. Such an adaptive control is performed by a fuzzy-logic controller, that takes into account both the throughput and frame retransmission count.

The paper [CT2] shows the performance of the IEEE 802.15.4 cluster-tree topology under large-scale RT WSNs scenarios. Simulation results in terms of throughput and delay are presented. Then, the key-aspects that limit the suitability for such kind of networks are discussed and viable solutions to the major problems are envisaged.

Another interesting aspect that was dealt with in this activity is the coexistence of multiple co-located IEEE 802.15.4 industrial networks, and particular attention was paid to cross-channel interference. The paper [CT3] summarises the results obtained on cross-channel interference in IEEE 802.15.4 networks.

-- Relaxing task isolation in soft real-time periodic systems (U. of Catania) --

This activity, continued from last year, showed that by "relaxing" task isolation, it is possible to efficiently deal with overruns in soft real-time systems with highly variable task execution times. The results obtained by the Randomized Dropping (RD) in [CT4] show that it is possible to bound task overruns in a probabilistic manner, thus providing "soft" task isolation while still maintaining system analyzability

-- Stochastic response time analysis of hybrid task sets in priority-driven soft real-time systems (U. of Catania) --

This activity, continued from last year, focuses on a novel task model, where a task is characterised by an Arrival Profile (AP) and an Execution Time Profile (ETP), both given by random variables with known distributions. In [CT5] results are given on the calculation of stochastic Response Time Profiles (RTPs) of tasks hierarchically scheduled using server-based techniques in a stochastic analysis framework.

-- Topology Management Protocols with Bounded Delay for Wireless Sensor Networks (U. of Catania) --

In Wireless Sensor Networks (WSNs) used in monitoring applications, the need to provide real-time traffic with an appropriate QoS typically clashes with the energy consumption constraints of the nodes, which have to work for long periods without the possibility of replacing their batteries. When both energy and QoS constraints are present, the role of topology control mechanisms is fundamental. This activity dealt with a topology control protocol to support energy-efficient real-time communication over WSNs. The aim of this work was to provide bounded delay for data traffic while reducing the energy consumption of the nodes. In [CT6] a detailed description of the protocol is given and both analytical and simulation results are presented that show the effects of the proposed topology control protocol in terms of reduced energy consumption, increased network capacity and reduced packet loss rate.

-- Architecture for dynamic service composition and algorithms for QoS-based composition of services (UC3M, Aveiro) --

The group at UC3M (affiliated to UPM) explored using the service-oriented paradigm to develop distributed real-time applications and, together with Aveiro, devised an architecture to support dynamic service composition. This kind of architecture can also be an alternative to provide some level of flexible resource management by supporting different profiles of each service, with different QoS and resource usage, and allowing a dynamic recomposition of the services/profiles involved in an application according to a predefined objective, e.g., maximize QoS of an application given the currently available resources, minimize the resources needed by a set of applications, etc. The architecture is based on a global entity, called the composer, which, together with a QoS manager, decides when and how to recompose a given application. The real-time coordination of the composition changes in the distributed system is carried out with the FTT-SE protocol.

Also, UC3M group has developed a series of algorithms for composition of real-time applications that are based on services. Applications are represented as graphs, where the nodes are the functionality pieces and the interactions are messages. In these algorithms, we find suboptimal solutions as a means to trade-off composition time with the finding of a solution. The search on the graph is based on the QoS parameters that characterise the services. This work has been done in collaboration the the U. Aveiro. [R10]

-- Distributed real-time middleware based on RTSJ – DRTSJ (UC3M) --

UC3M group has developed, in the starting phase of this project, extensions to RTSJ as a means to introduce more predictability in it to support the implementation of the distributed version of it (Distributed RTSJ – DRTSJ). Previous patterns as the NoHeapRemoteObject, Memory Area Pools, AGCMemory, threading model, etc., have been improved for their implementation in our prototype of DRTSJ called DREQUIEMI. Also, improvements to the communications protocol based on JRMP have been included in the prototype. [UC3M1] [UC3M2]

-- Design of a quality of service manager (SSSA, TUKL, UPM, Cantabria, UC3M) --

HOLA-QoS is a framework for managing QoS and resources and it has been used in media processing which UPM and UC3m have developed jointly. It is implemented as a layered architecture, so that layers can be replaced, as far as the API is kept. The work that has been finished is to replaced the lower levels of HOLA-QoS with two kernels with resource management facilities (also called resource kernels): MARTE (Cantabria) and SHARK (Pisa). The integration has been validated with some test applications. The final part of this work is the integration of the higher levels of HOLA-QoS that is currently being redesigned. One result of this work is the possibility of experimenting with the adaptation techniques that these advanced resource kernels provide.

The objective of the in ISO/IEC 23004-1 standard (Multimedia middleware, M3W) is to allow applications to execute multimedia functions with a minimum knowledge of the middleware and to allow applications to trigger updates to the middleware to extend the middleware API. Part 3 defines a component model, so that it is possible to use third party software in a seamless way. UPM has cooperated in the definition of a set of facilities for supporting resource and QoS Management in such a standard. In addition, an implementation of the resource Management components has been developed, taking HOLA-QoS as the basis for this work. The work done by UPM in the ASSERT Project has also relied on some of these concepts Some publications on HOLA-QoS can be found at http://www.dit.upm.es/str .

-- Design of a quality of service manager (UPM, UC3M) --

Budget accountant (BACC) is a software component that is able to handle resource budgets. In particular, it allows to assign budgets to entities, account for their usage on run-time and detect budgets overrun. There was a previous version of this component that has recently adapted to the linux kernel version 2.6. This work has been the basis for providing resource Management (CPU) facilities for the jamVM Java Virtual Machine (JVM). This will allow to provide CPU budgets to java threads. This approach is very relevant when due to different reasons (legacy code, libraries availability, …) it is not possible to use a real-time JVM.

Real-Time Languages (Cluster Integration)
Led by Alan Burns (University of York)

Partner teams (leaders): Alan Burns – University of York (UK), Michael Gonzalez-Harbour – University of Cantabria (Spain), Juan Zamorano – UP Madrid (Spain), Miguel Pinho – Polytechnic Institute of Porto (Portugal), Sergio Yovine – VERIMAG

Affiliated teams (leaders): Marisol García-Valls, Universidad Carlos III de Madrid (Spain).

Overview: The Ada language is still in use in many application domains, in particular the safety-critical areas such as avionics and railway signalling. The definition of the language itself has gone through a number of versions; the latest being Ada 2005. In the first standard the support for real-time embedded systems was weak with the concurrency model having a number of limitations. The Ada 95 version was a considerable improvement and did include a well defined set of primitives for undertaking fixed priority (i.e. essentially static) scheduling for non-adaptive applications. The research community, including members of the ARTIST2 community, has been involved in defining new language features that could extend the applicability of Ada, especially to the adaptive (more dynamic) domain of applications. Many of these features have found themselves incorporated into Ada 2005 (again due to the efforts of ARTIST2 members).

Ada 2005 is now defined and has undergone international standardisation. Currently implementation of all the real-time features of the language is awaited. It is appropriate therefore to assess the languages' expressive power in terms of the ease with which it will support the programming of flexible real-time systems. Much of the expertise surrounding Ada now lies within Europe, it is therefore important to build upon this situation to ensure the continuation of this lead. This will involve work within Europe and participation in international events, particularly in the US.

Supporting real-time functionality via language constructs rather than OS calls eases the programmer's task when writing complex applications. ARTIST2 partners have been involved in a number of standardisation activities and in ongoing research into real-time language primitives and associated analysis techniques. ARTIST2 provided the framework for this broad set of activities to compare outcomes and to influence each other's research. One particular area in which ARTIST2 partners are involved is the effort surrounding the use of Java as the core language on which real-time abstractions are built.

Work in Year 2

The activity started towards the end of Year 2. Effort was focused on an initial study of Ada 2005 and its implementation. Also work has been done to plan a series of activities concerned with the development of higher level abstractions for Ada. Part of this work involves work with international collaborators. A meeting/workshop on Ada 2005 was held in York as was an international open workshop on SCOOP.

Ada 2005 has a number of facilities that could make the programming of adaptive real-time systems much more straightforward and therefore likely to be used in an industrial context. But many of these features are new and have not been tested – in the sense of being used in an integrated way to build high-level abstractions. Work has started on this verification, and will continue in the following year.

Work in Year 3

The main focus in year 3 was again the Ada language, run-time support for the new language features (in Ada 2005) were planned (and partially implemented) and a library of reusable templates was instigated. A number of meetings were held to design the basic structures for these templates and to identify the main abstractions that needed to be supported. Work was also undertaken on implementing Ravenscar, a subset of Ada that is aimed at the safety critical domain.

The other language featured in the third year was Java and the Real-Time Specification (extensions) for Java (RTSJ). Work was undertaken on the memory management features and on extensions aimed at distributed applications. Again a number of key meetings and workshops were organised.

Final Results

-- Production of a Survey of Programming Languages (York, plus a large number of other ARTIST partners and non-ARTIST individuals linked to the ARTIST community). --

The production of real-time and embedded systems involves the use of many different tools and techniques. As these systems become more software centric, programming languages employed in the production of this software are now of crucial importance. Any language has the dual role of enabling expression whilst at the same time limiting the framework of concept and abstractions within which that expressive power may be applied. If a language does not support a particular notion then programmers cannot apply it and may even be totally unaware of its existence. For real-time programming languages there are many such concepts that are supported to a greater or lesser extent in a range of languages that purports to be appropriate for the embedded systems domain. For example: time, clocks, concurrency, deadlines, events and signals, exceptions, periodicity, scheduling and predictability are all important notions that programmers may wish to address and which should therefore be available to them via the implementation languages that they can employ.

This survey considers over twenty programming languages. The short reports available on each language aim to introduce the main features of the language, provide the links to further sources of information, and give an indication of the current developments within the language. Some languages are mature, used widely and are the subject of rigorous standardisation procedures. Others are research languages, with a small user population and an informal definition. All languages covered in the survey are implementation languages in that they are supported by tools (typically compilers) which generate executable code for the designated hardware platform. To give a structure to the survey each language is placed in one of five classes: imperative, functional, synchronous, model-based, and platform-based. However this is a loose classification as some languages could easily be placed in more than one category, and imperative languages can usually be constrained to mimic the other styles.

The survey in on the public ARTIST web site (follow links from home page, dissemination and Contribution to Standards) – or use the following.
http://www.artist-embedded.org/artist/ARTIST-Survey-of-Programming.html

The following have made contributions to this survey. Alejandro Alonso, Pierre Boulet, Frédéric Boussinot, Christian Buckl, Alan Burns, Magnus Carlsson, Paul Caspi, Jorge Coelho, Henk Corporaal, Susan Davidson, Víctor Fernández, Sebastian Fischmeister, Oana Florescu, Andy Gill, Thierry Gautier, Marc Geilen, Sébastien Gérard, Paul Le Guernic, Nicolas Halbwachs, Michael González Harbour, Leandro Soares Indrusiak, Julia L. Lawall, Insup Lee, Per Lindgren, Pieter Mosterman, Gilles Muller, Johan Nordlander, Luís Miguel Pinho, Juan Antonio de la Puente, Bran Selic, Robert de Simone, Bjorn von Sydow, Jean-Pierre Talpin, François Terrier, Bart Theelen, Stavros Tripakis, Marisol Garcia Valls, Eugenio Villar, Jeroen Voeten, Andy Wellings, Reinhard Wilhelm, Victor Wolf and Sergio Yovine.

-- *Development of Real-Time Utilities for Ada 2005 (Cantabria, Porto, UP Madrid, York)* --

The work started in year 3 is continuing with the development of patterns for real-time utilities. One aspect that was recognised during the second year was the limitation of the Ada 2005 model in that it did not allow requeue via interfaces. A report of this drawback was forwarded to the Ada language team and as a result the definition of Ada 2005 has been changed. This change has now been implemented by the main compiler team and is available.

Progress has not been as extensive as indicated last year by the initial lack of support for one of the key features of Ada 2005, namely group budgets. The current situation is that the feature is now available (see below), but extensive experience in using the feature has not yet occurred. This will occur in the future.

Finally, to note, extensions of the utilities to include support for mode changes has been progressing and will be reviewed at the next IRTAW in 2009.

A new technology has been developed for developing distributed and real-time component-based applications in Ada. The technology uses the “interfaces” mechanism, recently added in Ada 2005 to support multiple inheritance. Tools have been implemented to support the technology by automatic code generation of the elements that are needed to provide communication among parts of the application that are deployed in different processing nodes. This communication takes into account the real-time requirements specified in the deployment and configuration plan of the application. The approach allows full control of real-time properties both in the processors and in the networks, while eliminating the need of using a middleware layer.

Two middleware implementations providing support for distributed applications in Ada have been extended to support real-time requirements, both in the processing nodes and in the networks. The first of these implementations, called PolyOrb, supports different distribution technologies; among them, CORBA, and the Distributed Systems Annex (DSA) of the Ada specification. The second implementation is called Glade and uses the DSA technology. A comparison among the different implementations has been made providing information that is useful to decide under which circumstances one implementation offers advantages over the other one.

-- *Implementation of Ada 2005 (Cantabria)* --

AdaCore, a leading Ada compiler vendor that develops and distributes the gnat free software compiler, now distributes the MaRTE operating system developed at the University of Cantabria, as one of the run-time systems that provides underlying support to the concurrency of Ada tasks. This run-time system, available on Linux platforms, offers all of the features defined in the real-time Annex of Ada 95, and some of the features defined in Ada 2005. In the current reporting period work has continued to include in this implementation the remaining Ada 2005 services. Implementation of task group budgets has now been completed. In

addition, implementation efforts have been made to better integrate MaRTE OS with the gnat run-time system.

-- Ada 2005 real-time mechanisms (UP Madrid, York) --

The UPM team have been working on implementing the Ada 2005 real-time mechanisms on the Open Ravenscar real-time Kernel (ORK), aimed at high-integrity systems using the Ravenscar profile of Ada tasking on SPARC 8 (LEON) computers. A new version of the kernel, ORK+, has been developed, which includes execution-time clocks and timers, timing events and group budgets. Since execution-time timers and group budgets are not allowed in the Ada 2005 definition of the Ravenscar profile, an extended profile has been defined that includes up to one execution-time timer per task, declared at library level. ORK+ is a core component of the ASSERT Virtual Machine, the execution platform for real-time components developed in the ASSERT project.

The kernel is integrated with the GNAT GPL 2007 compilation system, and can be downloaded from http://www.dit.upm.es/ork

-- Real-Time Java Community Building (UC3M, York) --

The real-time Java user (academic and industrial) community is still fragmented. However, the series of Real-Time Java Workshop (JTRES – Java Technology for Real-time and Embedded Systems), is now firmly established. The 2007 workshop was held in Vienna and the 2008 one at Sun Microsystems Santa Clara campus. York has played a leading role on the steering committee for these workshops.

Within the ARTIST partners, the main Java-related research work is being done by UC3M and York. UC3M is working on the distributed version of RTSJ; from the implementation of a real-time RMI, the group has identified, specified, and implemented over the past years a number of extensions to RTSJ. These extensions are a means of facilitating the implementation of real-time distributed applications based on RTSJ by transferring to the middleware (DRTSJ) the job of guaranteeing real-time interaction. York is addressing a range of issues including asynchronous event handling, mobile code, multiprocessor issues.

-- Maintaining Momentum: International Standards Work (York) -

The main implementation approach of real-time Java, the Real-Time Specification for Java (RTSJ) continues to evolve. Over the last year, the main work has been undertaken in two Java Specification Requests (JSR 282 and JSR 302).

JSR 282 is, among other things, trying to fill some of the gaps that were left in the original specification, and to begin to explore more direct supports for SMPs. Proposed enhancements range from changes in scheduling models to improved memory management issues to offer better support to its region model. This process is going to be useful not only for programmer but for other specifications which will act as the acid test for RTSJ. JSR 302 is considering safety and mission critical profiles. York plays a major role in both these activities.

Distributed real-time Java technologies require an extraordinary effort in order to integrate the current centralized real-time Java languages with the traditional middleware distribution paradigms. JSR 50 leads this initiative; however, currently, there the work on the distributed version of RTSJ is stopped and no newer drafts have been released since 2006. Considering that distribution will become a key in the success of Java-based distributed applications with real-time requirements, the work carried out at UC3M is aimed at the construction of an architecture which offers platform independence (see below).

-- Maintaining Momentum: Research Work (UC3M, Verimag, York) --

This UC3M group has focused on Distributed Real-Time Java Issues. The work provides a set of important abstractions such as support for a distributed garbage collection and a naming service, more related to the RMI (Remote Method Invocation) model. Also, it contains other contributions like the possibility of using asynchronous remote invocations or centralized synchronization services which are useful in the development of many real-time systems. An initial programming interface, called DREQUIEMI, for this computation model has been already defined; it is aligned with the RMI (Remote Method Invocation) and the RTSJ (Real-time Specification for Java) technologies. Besides, in the specific context of the RTSJ, a set of extensions (to the region model, to the reference model and to the threading model) has been proposed to simplify the development of both, centralized and distributed real-time applications. In the course of the last year, the group of Carlos III University of Madrid has carried out suite calculations and measurements to determine the drawbacks in communication when using the complex layer structure of Sun's Java RMI. This practical side identifies room for improvements in the remote invocation theoretical model, current implementation, and mismatches on the integration of RMI and RTSJ to become DRTSJ.

From the perspective of memory management, the results of UC3M on the implementation of parts of DRTSJ, especially on memory management issues, link the work of the group with JCP-282. UC3M has worked on the evaluation of current programming extensions designed to violate assignment and single-parent rules that are insufficient for the demands of a distribution middleware. Current violation mechanism, named portals, cannot be accepted as valid mechanism to gain access to an arbitrary region; an extension to the current programming interface is required in order to offer more general facilities which make such a process easier. Furthermore, this enhancement has to extend not only the weak semantics of Java (now considered as an item in the JCP-282 list), but also the case of strong references, which now are forgotten by the JCP-282. The second is the need of other region models which offer improvements to the current region model, thus giving some kind of extra support for specific demands. For instance, current infrastructures do not allow to recover regions automatically once they are empty, one mechanism which may be useful to generate a valuable support for no-heap distributed remote objects in DRTSJ. UC3M has also further elaborated the concurrency model of RTSJ. The group's experience in DRTSJ implementation demonstrated the usefulness of having a more flexible threading infrastructure for RTSJ.

Java-based distributed real-time systems require appropriate memory management. Since currently no 100% real-time garbage collection is possible, predictable memory management relies mostly on its avoidance in favor of using memory regions. Currently, RTSJ offers a region-based model for the development of centralized systems, but its distributed equivalents lack a similar support. The group of Carlos III University of Madrid has developed the "no-heap remote object" paradigm that allows to remove the objects created during a remote invocation automatically. The model extraordinarily reduces the number of changes required in the platform to support it. On the other hand, VERIMAG worked on the direction of trying to provide predictable dynamic memory management. For this, the group developed a static analysis capable of providing a safe but tight approximation of the peak memory occupancy required to run a program using RTSJ scoped-memory allocation.

The EU JEOPARD project has recently been started with a remit to investigate SMP and multicore issues. York is a member of the team. In order to make the RTSJ fully defined for SMP multiprocessor systems, the following issues have been addressed.

- The dispatching model -- the current specification has a conceptual model which assumes a single run queue per priority level.

- The allocation model -- the current specification provides no mechanisms to support processor affinity,
- The synchronization model -- the current specification does not distinguish between synchrononized methods which suspend holding their locks and those that do not (this is important for multiprocessor synchronization algorithms),
- The cost enforcement model -- the current specification does not consider the fact that processing groups can contain scheduling objects which might be simultaneously executing,
- The affinity of interrupts (happenings) -- the current specification provides no mechanism to tie interrupts (happenings) and their handlers to particular processors, and
- The failure model -- the current specification makes no statements about partial failures of the underlying platform.

Dynamic and Pervasive Networking (Cluster Integration)
Led by Eduardo Tovar (Porto)

Partner teams (leaders): Eduardo Tovar – Polytechnic Institute of Porto (Portugal), Luís Almeida – University of Aveiro (Portugal), Giorgio Buttazzo – University of Pisa (Italy), Alan Burns – University of York (UK).

Affiliated teams (leaders): Lucia Lo Bello – University of Catania (Italy), Pau Martí – Polytechnic University of Catalonia (Spain), Marisol García-Valls – Universidad Carlos III de Madrid (Spain), Julián Proenza – University of Balearic Islands (Spain), Wilfried Elmenreich – Technical University of Vienna (Austria), José Maria Giron – Complutense University of Madrid (Spain).

Overview: Looking at the current scenario in embedded systems we see a consistently growing role of networking, ranging from the interconnection of autonomous devices such as cellular phones, personal digital assistants (PDAs), laptops and their peripherals, to the provision of pervasive access to multimedia and telecommunication networks, to the deployment and operation of large-scale sensor networks, to intelligence distribution in complex embedded systems, or even, at a small physical scale, to connect multiple processing cores within Systems-on-Chip (SoCs).

In this vast horizon, the activity on Dynamic and Pervasive Networks of the ARTIST2 Adaptive Real-Time Systems (ART) cluster focuses on Wireless Sensor Networks (WSNs) and Networked Embedded Systems (NESs), areas in which many open challenges remain, and many new problems arose in the past few years. From energy-aware communication to efficient data aggregation, real-time routing and operation under severe resource constraints in WSN, to dynamic topology tracking and management and dynamic team composition in MANETs, to Quality-of-Service (QoS) adaptation and higher software integration in NES, these are all open issues in which further advances will impact the embedded systems community at large, from the way systems are designed to the way they interact with users and the way they are immersed in the environment.

Work in Year 3

This activity started in Year 3.

Year 4 results are a natural continuation of the effort carried out during the first year (Year 3) of the DPN activity, so most of the research frameworks span over these two years.

In Year 3, we devoted main research efforts to aspects related to *wireless sensor networks*, to *re-configurability and on-line adaptation* mechanisms in networked embedded systems and to the use and improvements in standard and COTS communication technologies for WSNs, MANETs and NESs. A substantial part of this work was based on the IEEE 802.3/Ethernet, IEEE 802.11/WiFi and IEEE 802.15.4/ZigBee wireless protocol standards, which were analysed, improved, engineered, implemented and tested to be used in time-sensitive applications.

Final Results

a) Worst-case analysis and dimensioning of cluster-tree wireless sensor networks

Modeling the fundamental performance limits of Wireless Sensor Networks (WSNs) is of paramount importance to understand their behavior under worst-case conditions and to make the appropriate design choices. This is of particular importance for time-sensitive WSN applications, where the timing behavior of the operating system (task execution must respect deadlines) and of the network protocols (message transmission must respect deadlines) impacts on the correct operation of these applications. In that direction, researchers from **Porto and Prague** have proposed a methodology for modeling cluster-tree WSNs where the sink can either be static or mobile [PO14, PO15]. This methodology enables the computation of the worst-case end-to-end delays, buffering and bandwidth requirements across any source-destination path in a cluster-tree WSN. This generic methodology was instantiated for the particular case of IEEE 802.15.4/ZigBee cluster-tree WSNs and validated through a comprehensive experimental study using commercially available technology, namely TelosB motes running TinyOS.

Catania addressed the limitations of the cluster tree topology [CT1], namely, low scalability and unbalanced energy consumption through simulation results obtained using the Omnet++ tool, and proposes viable alternatives to address these problems.

b) QoS add-ons to the IEEE 802.15.4 and ZigBee protocols

Porto has continued the ART-WiSe research framework around the use of IEEE 802.15.4 and ZigBee as federating communication protocols for WSN applications with QoS requirements (energy-efficiency, timeliness, throughput, reliability). In this context in this reporting period, several mechanisms were proposed in order to improve QoS, such as a Hidden-Node Avoidance Mechanism for wireless sensor networks [PO1], a Time Division Beacon Scheduling mechanism for ZigBee cluster-tree networks [PO2], a mechanism for improving bandwidth usage of Guaranteed Time Slot (GTS) /in IEEE 802.15.4 [PO3] and the outline of a gateway between IEEE 802.15.4/ZigBee and IEEE 802.11/WiFi [PO08].

Notably, the open-ZB open-source toolset (http://www.open-ZB.net) for the IEEE 802.15.4 and ZigBee protocols [PO5, PO9, PO16] witnessed an outstanding number of visits (over 53000) and downloads (over 3000) since December 2006.

Another issue that has been addressed by **Catania** is cross-channel interference in co-located IEEE 802.15.4 industrial networks [CT2]. The problem is tackled from two different perspectives and provides both analytical and experimental results. The latter were obtained through an extensive series of measurements run in order to assess the performance of IEEE 802.15.4 networks under different critical operating conditions. The analytical results are based on the properties of the coding used at the Physical layer of the IEEE 802.15.4 protocol, in particular the power spectral density of the signal.

c) Supporting real-time communications in wireless sensor networks over the ERIKA real-time operating system

IEEE 802.15.4/ZigBee and TinyOS have been playing an important role in leveraging a new generation of large-scale networked embedded systems. However, based on previous experience (from **Porto** and **Prague**) on the implementation and use of the IEEE 802.15.4/ZigBee protocols over TinyOS (http://www.open-ZB.net), several problems (producing loss of synchronization and even network crashes) emerge due to some limitations of TinyOS, namely related to the lack of task pre-emption and prioritization [PO10]. This unreliability is not a major concern for non-critical environments where the nodes are supposed to guarantee best-effort services. However, when real-time guarantees are required, different software solutions must be used to support real-time services in such networked applications. In this context, researchers from **Porto** and **Pisa** have been working [PO17] towards the implementation of the IEEE 802.15.4/ZigBee protocol stack over ERIKA, a real-time operating system for resourceconstrained embedded systems.

d) Keeping low the time-complexity of distributed computations of physical quantities in large-scale and very dense sensor systems.

WiDom (Wireless Dominance protocol) and WiSe-CAN (Wireless Sensor Networks protocol based on the Controller Area Network protocol) are two related research efforts that have evolved through this year involving researchers from **Porto** and **Vienna**. We have (i) shown how to deal with sensor faults [PO11], (ii) shown how the approach can be used to perform localization [PO12] and (iii) how the approach can be used with the CAN bus [PO18].

e) Robust communication with star topologies

When considering wired communication, star topologies present several advantages in robustness with respect to other topologies such as the typical buses, e.g., in terms of error confinement. This lead to two different developments, one with CAN networks carried out by **Mallorca** and **Aveiro**, and another with switched Ethernet, by **Aveiro**, with collaboration from **CMU**. In the former case, two star topologies were developed for CAN, namely CANcentrate, which is a simplex star, and ReCANcentrate, which is replicated. As a consequence of this work a book chapter [AVc1] was produced as well as a paper that has been accepted for publication in IEEE Computer [AVc2]. This year in particular, special attention was devoted to data consistency issues in ReCANcentrate, during dynamic phenomena associated to transient errors [AVc3] [AVc4].

Concerning switched Ethernet, current COTS switches do not include protection against faults in the time domain, or even overloads. This is particularly important in dynamic systems in which the current configuration can change at run-time. However, to provide such policing and protection, the switch must know the properties of the streams currently crossing it. Therefore, the team at **Aveiro**, with the collaboration from **CMU**, designed a new switch that carries out traffic scheduling with resource reservation, controlling the transmissions, and verifying compliance of the incoming streams with their negociated properties at the input. Non-compliant packets can be promptly eliminated, thus not interfering with the remaining system. A preliminary version of such switch is presented in [AV1].

f) Schedulability analysis for specific traffic types

The transmission of long messages in CAN requires fragmentation but the medium access control still operates on a packet basis. In some applications it was found that it could be convenient to transfer such messages at once, overriding the packet level MAC. A specific schedulability analysis was developed by **Pisa** and **Aveiro** for such cases [PIc1]

g) real-time support to middleware layers

Providing real-time support to existing middlewares requires appropriate infrastructures. In this case, **Aveiro** has provided the FTT-SE platform that has been used in **Madrid** (UC3M) to support composition of service-based applications [AVc5]. **Cantabria** has provided the FRESCOR contracting middleware and Aveiro supported it on FTT-SE to allow centralized contract management and facilitate contract run-time adaptation [AVc6]. Finally, Madrid (UC3M) has also developed different enhancements for RTSJ, especially in memory management [MA1] [MA2] that introduce higher predictability and efficiency towards implementing the Distributed RTSJ. These extensions have been implemented in their DRTSJ prototype named DREQUIEMI.

h) Educational Testbed supporting teaching of industrial wired/wireless networks

In order to enable students to understand how to properly design hybrid wired/wireless industrial communication networks, the adoption of new concepts and paradigms is needed. The work [CT3] describes the design and implementation of an educational testbed for a course in industrial communication, via a remotely accessible platform emulating various network configurations based on user configured network conditions, also providing the user with monitoring capability.

Scheduling inside the routers is one main focus of the testbed. When the network is being configured, students can choose the scheduling policy they want to investigate and even assign the various routers different policies in order to evaluate the effect of each of them on the router behaviour. In addition to the First-In–First-Out (FIFO) policy, which is present in all traditional routers, the testbed also uses some scheduling techniques which adapt the Earliest Deadline First (EDF) algorithm for use in a packet-switched multihop network.

2.2.3 Compilers and Timing Analysis Cluster

This cluster is composed of the following activities:

Timing Analysis (Platform)
Led by Reinhard Wilhelm (Saarland University)

Partner teams (leaders): Björn Lisper – Mälardalen university (Sweden), Reinhard Wilhelm – Saarland University (Germany), Christian Ferdinand – AbsInt GmbH (Germany), Guillem Bernat – University of York (UK), Jan Gustafsson – Mälardalen university (Sweden), Andreas Ermedahl – Mälardalen university (Sweden), Peter Puschner – TU Vienna (Austria), Niklas Holsti – Tidorum Ltd. (Finland), Peter Marwedel – TU Dortmund (Germany), Stephan Thesing – Saarland University (Germany), Raimund Kirner – TU Vienna (Austria) , Oleg Parshin – Saarland University, Saarbrücken (Germany), Jan Reineke – Saarland University, Saarbrücken (Germany), Sebastian Altmeyer – Saarland University, Saarbrücken (Germany), Gernot Gebhard – Saarland University, Saarbrücken (Germany).

Affiliated teams (leaders): Isabelle Puaut – IRISA (France)

Overview: Europe is leading this field. The only commercially available WCET tools, aiT, Bound-T, and RapiTime, and most of the academic prototypes, Heptane, SymTA/P, SWEET, and the tools of TU Vienna, are of European origin.

The commercial tools and the academic prototypes both follow two different approaches, namely analytical and measurement-based approaches. Each individual approach has to solve essentially the same set of sub-problems. The differences between the approaches lie in the methods used to solve some of the subproblems. Methods for some subproblems can be combined across approaches.

The timing-analysis problem has been solved for tasks executed on monoprocessors without preemption. Projects for the development of time-critical systems in industry use tools of ARTIST2 partners. However, there is room for improvement. The usability can be improved by reducing the necessary amount of user interaction. The efficiency of the tools can be improved by improving our knowledge about the underlying processor architecture. The tool realization can be simplified by automated generation from formal descriptions of the processors and systems to be analysed. The different tools offer different strengths in the areas of different subproblems. The goal is to combine the respective strengths.

The transition of industry to multi-processor and multi-core targets opens up a whole domain of new research problems. Not all envisioned multi-core architectures will allow timing analysis. The timing-analysis community needs to take an influence on the architecture design to make timing analysis feasible. Good compromises have to be found between the goals of good average-case performance, good predictability of the worst case and the necessary effort for the analyses.

Proposals for a modular framework allowing the integration of prototypical developments existed before the start of the NoE. Documented and supported interfaces existed for the individual tools, but not for any combination of components from different tools.

In an ongoing discussion between the participants about WCET-tool architectures, interfaces, and integration agreement on interfaces has been reached.

Traditionally, timing analysis tools have been designed independently of compilers. It has now turned out that proceeding along this path would result in a duplication of efforts. Flow facts are available in compilers and need to be regenerated in timing analysis tools. Timing information is available in timing analysis tools and would be useful for timing-aware optimizations in

compilers. Currently, compilers use very rough approximations of timing, if they use any timing model at all. As a result, the impact of certain transformations on the run-time is frequently not known by the compilers. Hence, the user has to follow a trial-and-error approach, experimenting with different compiler options and figuring out a suitable combination of them. However, even this time-consuming process cannot really minimize the execution time since options which might be good for some part of the code might lead to bad result for some other part of the code. A tight integration of timing models into compilers and their optimizations is urgently needed.

Work in Year 1

Work Achieved in the First 6 Months

We started with an extensive discussion of potential tool-interface languages. Such a language needs to offer the representation of object programs with an adequate attribute mechanism to store annotations about feasible and infeasible paths for program execution and to present analysis results at the program-source level in sufficient detail to be easily understandable by the tool user.

The prime candidate for such a language was CRL2, the interface language of AbsInt's timing-analysis tool aiT. CRL2 is the result of long evolution of intermediate representations in analysis frameworks of AbsInt and Saarland University.

The commercially-supported CRL2 format was interfaced with several of the components needed for the planned Timing-Analysis platform, i.e., parts of the analysis tool chain made ready to communicate via CRL2. With the selection of CRL2 a robust and generic interface was introduced into the ARTIST2 framework.

CLR2 is a generic and processor independent format usable for static analysis (including WCET analysis), optimisation of machine code and assembly language. It supports the integrated representation of control flow graph and intermediate analysis results. An efficient C/C++ library reads/writes CRL2 interface files in a text-representation format and provides an API to the data structures used by the components of the timing-analysis tool suite.

CRL2 has an interface to the Program-Analyzer Generator developed at Saarland University, such that interprocedural analyses are easily implemented.

For the given reasons, CRL2 was first chosen for the integration of various work groups' analyses. Several partners in the Timing-Analysis cluster started to interface with CRL2. Experiments at IRISA and Tidorum with the CRL2 library of AbsInt showed good results. In particular, preliminary experiments made at IRISA used CRL2 to add a tree structure on top of the control flow graph supported CRL2 library. The experiments showed that the CRL2 library, through its attribute system, is flexible enough to define data structures used by WCET analysis methods different from those implemented in aiT.

Based on the first successful results with CRL2 it was discussed that the interchange format should not only be used for timing analysis but should also serve as a compiler-analyser interface, thus facilitating the WCET-aware compilation of code. The interface language is called AIR, for ARTIST2 Intermediate Representation. AIR is based on CRL2. CRL2 is a dialect (superset) of AIR as explained in Section 2.2. Further, it was found that in order to suit the needs of the partners of the CTA cluster, a text format of the interchange format had to be defined and the documentation should be extended (until then CRL2 was available as a library only). AbsInt offered to provide this format definition and documentation.

Besides the work on the common interchange format for timing analysis, the partners of the CTA cluster accomplished the following:

- Students from Mälardalen University performed a number of industrial WCET analysis case studies in Swedish enterprises, mostly with AbsInt's tool aiT. The case studies showed that for common embedded processors tight WCET bounds could be obtained. However, the effort needed to provide the required program-flow information by hand turned out to be high and time consuming.
- Mälardalen and Vienna developed a prototype editor plug-in that helps the programmer to achieve more predictable timing and aims at simplifying WCET analysis. By highlighting code segments that are executed conditionally, i.e., that are executed only for a subset of all possible inputs to a piece of code, the editor helps the programmer identify possible sources of timing variability. Avoiding such input-dependent conditionals, or at least minimising them in length, leads to code that in general has fewer execution paths, a smaller execution-time jitter, and is thus easier to analyse for its WCET.

Work Achieved in Months 6-12

The major focus of the activity was on the further development of the common interchange format for WCET analysis:

We identified one important area that was not explicitly represented in CRL2: the "computation semantics", by which we mean the computation that is performed by each instruction in the program. As CRL2 has been used so far, the control-flow graph identifies the instructions and their operands, but an analysis tool must itself have enough knowledge of the target processor to map the instruction identifier and operand identifiers to the computation, for example to understand that the instruction adds two registers and puts the sum in a third register. This computation, connected to the control-flow graph, is the essential information that the tools from Mälardalen and Tidorum use to find constraints on execution flow, such as loop bounds. We therefore decided to define an explicit and general representation of this computation semantics as one part of AIR, using the flexible attribute mechanism from CRL2. Work on this extension is under way.

Peter Marwedel's group at Dortmund University used CRL2 to interface their compiler with the aiT tool. Similarly to earlier activities, one of the results was that a better documentation on aiT's usage of attributes would be useful.

While AbsInt and Saarland University were working on extensions and improvements of CRL2, feedback from Rennes, Mälardalen and Tidorum asked for a formalisation of the control flow graph structure and some further extensions of the interchange format.

In parallel, Vienna and Mälardalen started work on extending the path-annotation support of CRL2. As for the computation semantics, the plan was to use the attribute mechanism of CRL2 to represent these annotations about feasible and infeasible execution paths to facilitate a highly accurate modelling of the possible execution paths for WCET analysis. Inputs from Vienna, Mälardalen, Tidorum, and AbsInt for these path annotations have been collected and summarized in a presentation.

Though each of the activities was very promising by itself, the consortium had to find that the definition of such a complex interface as the WCET interchange format requires and would yet require a lot of further work till its final completion.

Besides the core work on the common format the partners of this activity reported about the following work related to timing analysis:

- Tidorum and Mälardalen cooperated on implementing a version of the Bound-T WCET analysis tool for the Renesas H8/300 processor, which can be found in the popular Lego Mindstorms kit. Mälardalen is now using this tool in real-time systems education.

- Licensing policies and the issue of securing the availability of AbsInt's CRL2 library had been discussed. We are looking for a full and open definition of the syntax and semantics of AIR so that anyone can build their own AIR libraries.
- Vienna and York started to cooperate on measurement-based timing analysis. The partners had been working individually on measurement-based analysis before. York uses measurement-based WCET analysis for non safety-critical applications and has a strong focus on code instrumentation and report formats for representing measurement details to the user. In Vienna, the use of measurements is seen as a complement to static analysis for the purpose of validating static-analysis results. In Vienna, the focus is on the automatic test-data generation. The cooperation started within the CTA cluster aims at combining the efforts and exchanging the complementary know-how of the groups. In particular, the partners started to work on the definition of coverage criteria for measurement-based WCET analysis.
- The industrial WCET case studies performed by students from Mälardalen University in Swedish enterprises were continued. The findings essentially confirmed the conclusions from the earlier studies.
- Representatives of AbsInt, Tidorum and Mälardalen all demonstrated their WCET tools and participated in the Real-Time in Sweden (RTiS) conference, held in Skövde, Sweden, Aug 2005. This was a joint effort to introduce the concept of static timing analysis to the Swedish companies that participated in the RTiS conference. See: http://www.snart.org for details about RTiS.

A number of publications have been produced as a result of the cooperation inside the cluster. They are included in the list of publications collected and submitted by the coordinator.

Work in Year 2

Definition of AIR (ARTIST2 Intermediate program representation for WCET tools)

In the past few months, a file format specification was developed to define the AIR ('ARTIST2 Intermediate') format that may be used by the cluster participants' tools for integration. So AIR is the proposed exchange format of the tools of the groups participating in the ARTIST2 project. The format is based on CRL2, which is the successor of CRL. These formats were originally developed in cooperation by Saarland University and AbsInt Angewandte Informatik GmbH over several years of work.

The idea behind AIR is that an interface is to be defined on the file-format level, in contrast to CRL2, whose interface definition only covers the C++ library interface. Internally in AbsInt tools, a specification of the C++ library interface is preferred over a file format specification, simply because all tools use the library and thus the storage on disk is secondary. For ARTIST2, different work groups prefer their own libraries over the usage of proprietary software, so there is a serious demand for a file format specification.

Since CRL2 was not primarily meant to be a file format, much work had to be done before this document could be written. Apart from the mere documentation the file format had to be defined and implemented. In order to get a stable interface on file level CRL2 had to be extended. For example, version numbers and specification IDs had to be added to meet the strict safety criteria of real-time systems analysis. Thus, this document can be viewed as the first step of the final documentation phase in a larger effort towards an exchange file-format for the different WCET tools and tool components used within this ARTIST2 activity.

From the release of the first AIR specification on, the CRL2's file format interface will be a dialect of the AIR file format. CRL2 as well as dialects of other work groups are allowed to

feature extensions as long as they are not vital for the operation of the tools. E.g., AbsInt tools will only use the plain AIR file format during normal operation, the extensions of CRL2 are mainly implemented for debugging and diagnosis purposes. In the same way, extensions of other dialects shall never be vital to the operation of the corresponding tools
http://www.absint.com/artist2/doc/crl2/air.pdf

Timing-Analysis Survey Paper

The work on the survey paper about Timing-Analysis Methods and Tools has clarified the characteristics, the advantages and disadvantages, and the application domains for the different approaches, e.g. analytical and measurement-based approaches. It has clarified the modularisation of the overall timing-analysis task and the possibility of combining modules for different subtasks across the approaches. The joint authorship of this paper expresses strong and enduring cooperation between the different groups. This paper will represent a landmark publication for the area!

Timing Anomalies Characterisation and Checking

Timing Anomalies in processors produce counter-intuitive timing behaviour, i.e., local worst-case behaviour does not necessarily lead to global worst-case behaviour. The existence of timing anomalies requires complex timing-analysis procedures. Saarland University together with Freiburg University have worked on the clarification of the concept and the origins of timing anomalies. The goal is an automatically checkable definition of timing anomalies, which would allow for a safe reduction of the WCET analysis effort whenever the absence of timing anomalies can be shown for a processor platform. Furthermore, it was found that certain cache-replacement strategies lead to Domino Effects, which are timing anomalies without bounds for their effects. This work is funded by the Transregional Research Centre AVACS (Automatic Verification and Analysis of Complex Systems) of the Deutsche Forschungsgemeinschaft.

Parametric WCET Analysis

The runtime of programs might depend on parameters. In these cases the worst case execution time (WCET) has to be recomputed for each parameter assignment. This can be very time consuming. On the other hand the relation between parameters and WCET cannot be easily identified.

Saarland University together with Mälardalen University have initiated collaboration about this type of parametric WCET analysis. Two M.Sc. students from Saarland visited Mälardalen in early 2006 to learn about the Mälardalen approach.

In the joint approach a parametric WCET analysis based on the aiT-tool chain is performed. It computes a WCET formula instead of a concrete value. Since programs spend most of their runtime in loops, we focus on a parametric loop-bound analysis. Prior to this part the parameters of the executable have to be determined. In the path analysis part, a parametric optimisation method is needed. Afterwards the resulting formula has to be evaluated. Evaluation in this context means visualisation or instantiation of the formula.

WCET Analysis Benchmark Suite

Mälardalen University has collected a suite of benchmark programs for WCET analysis. The suite is maintained on the web by Mälardalen. One of the purposes of this benchmark suite is to be able to evaluate and compare different WCET tools as is done in the initiated WCET Tool Challenge (cf. http://www.mrtc.mdh.se/projects/wcet/benchmarks.html

Input Format for Flow Analysis

Mälardalen University has initiated work to define a standard input code format "ALF" for flow analysis. The purpose is to facilitate flow analysis of codes in different formats by translating to ALF. ALF will provide an interface to Mälardalen's flow analysis. A first draft for ALF exists, and it will be disseminated within the cluster before the format is finalised. ALF should also be harmonised with the AIR instruction semantics format.

Synergy between Code Synthesis and Timing Analysis

Saarland University together with AbsInt and ETAS have integrated the ASCET-SD code-synthesis with AbsInt's timing-analysis tool to improve usability of the timing-analysis tool and precision of the results.

Timing Predictability

First quantitative results have been obtained on the influence of architectural properties on the timing predictability of embedded systems. In particular, four different cache replacement policies and their influence on predictability have been considered at Saarland University. This research is funded by AVACS. On the side of TU Vienna, a model for a time-predictable processing node has been worked out – on this node software timing behaviour can be predicted with the granularity of the CPU clock. This node uses a purely time-triggered input-output interface and relies on single-path code (code that is free from input-data dependent control flow) in both the operating system and the application code. Tasks are only preempted at pre-planned task preemption points and simple clock synchronization keeps the operations of the nodes in synchrony with its real-time environment. The work on the time-predictable node yielded a time-predictable task-preemption model where an instruction counter instead of the CPU clock is used to implement preemptions (It was shown that CPU-clock based pre-emption may lead to unpredictable timing).

Measurement-Based WCET Analysis

As the complexity of WCET analysis varies with the structure of the program to be analysed and the type of target hardware, TU Vienna and York worked out a detailed list of issues for measurement-based WCET analysis. This list of issues is to be used for different purposes: First, it is a check list for the designer of a WCET analysis tool. Second, it gives the system developer clues about relevant hardware and software criteria when designing a system with the goal of simple analysability and predictability. The list is divided into three categories: a) issues that only relate to the software of the system, b) issues that address only the target hardware of the system, and c) issues that are relevant for both, the software and the hardware part of the system. The partners used these results as a starting point for the work on coverage criteria for measurement-based WCET analysis that was initiated in this work period. The goal of this work item is to find meaningful metrics for assessing the timing-related code coverage and the value of input-data sets for the measurement-based analysis.
The results are documented in a technical report, which is available at:
http://www.vmars.tuwien.ac.at/php/pserver/extern/docdetail.php?DID=1975&viewmode=published&year=2007

Work in Year 3

Timing Analysis and Timing Predictability (USaar and AbsInt)

Timing Anomalies are a difficult problem for timing analysis. They complicate the design of tools; they increase the necessary analysis effort, and they decrease the precision of the

results. No real understanding of the concept existed, only observational properties were known. A new definition has been produced that helps to understand this phenomenon. It covers both scheduling as well as speculation anomalies.

The large state space to be traversed during the pipeline analysis largely determines the necessary analysis effort. An approach has been followed to use symbolic representations of this state space to increase efficiency of the analysis.

The work on predictability has continued and first hard analytical results about predictability of architectural features have been obtained, in this case cache replacement strategies. These show that the replacement strategy has a strong influence on the precision of any type of cache analysis.

The formal derivation of abstract processor timing models has been mostly implemented. This process starts from a specification of the hardware architecture in VHDL and proceeds by a series of analyses and transformations. Analyses of such models for several kinds of properties will be possible once formally derived abstract architectural models are available.

Preemptive scheduling of hard real-time tasks requires precise estimations of context-switch costs. These are largely dependent on the cache-refill costs caused by pre-empting tasks. An approach has been developed and implemented that estimates and even minimizes the cache interference of tasks. The latter optimization uses the memory allocation to define the cache mapping.

Synergy between Code Synthesis and Timing Analysis (USaar and AbsInt)

An integration of AbsInt's aiT timing-analysis tool with the ASCET specification and synthesis tool of ETAS has been realized, and experimental results about the effect have been produced (ISoLA 2006 Paper, see 2.3.4).

ARTIST Interchange Representation and Attribute Database: AIR (AbsInt)

The work on AIR has continued. The syntax specification produced in Year 2 is now accompanied by a data-base that defines the current set of "attributes" that carry much of the inputs and outputs of the analysis as decoration on the extended control-flow graph. A data base for managing these attributes was designed and implemented. The data base specifies attribute names and types, their location (e.g. at routines or at instructions) and their access rights (which tools may introduce/read/write which attributes). The attribute data base was established on a web server and is accessible by all ARTIST2 partners. The database interface is designed to be open to allow the definition of further attributes by other ARTIST participants.
http://www.theiling.de/absint/attrdb.fcgi

Computation Semantics Representation: ALF (USaar, Tidorum, Mälardalen, AbsInt)

To support the flow analysis (loop bounds and infeasible paths) AIR must be able to represent the computations executed by the instructions in the program under analysis. CRL2, the basis of AIR, does not have an explicit, portable computation semantics representation. Work on such a representation started in Year 2 and continued in Year 3. The most advanced candidate is the ALF language from Mälardalen University. Mälardalen produced a preliminary specification of the ALF syntax and semantics
http://www.mrtc.mdh.se/index.php?choice=publications&id=1351

WCET Challenge 2006 (Mälardalen, USaar, AbsInt, Tidorum, National University of Singapore, TU Vienna together with Christian-Albrechts University Kiel, DaimlerChrysler Research Ulm, University Duisburg-Essen, and University Stuttgart)

The purpose of the WCET Tool Challenge was to be able to study, compare and discuss the properties of different WCET tools and approaches, to define common metrics, and to enhance the existing WCET benchmarks. Four WCET tools (two commercial, two research prototypes) completed the Challenge. The WCET Tool Challenge was performed during the autumn of 2006 and resulted in two papers at ISoLA 2006. The WCET Tool Challenge was also presented at the WCET workshop 2007.
Based on our experience with the WCET Challenge 2006 we plan to hold future challenges bi-annually with results presented at the International WCET Workshop. The next challenge will be run in the fall of 2007 and spring of 2008, with a more formal and rigorous structure and hopefully more participants. The results will be presented at WCET-2008.
http://www.idt.mdh.se/personal/jgn/challenge/

Transformation of Flow Information during Compilation (TU Vienna/Real-Time Systems Group + TUV/Programming Languages Group)

Flow information has to be used in general to guide the worst-case execution time analysis. When given such extra flow information manually, it is most convenient to give them at source code level. However, when compiling the code the compiler may change the control-flow structure of the program, thus rendering the original flow information as invalid. Thus, it is necessary to update the flow information in parallel to the compiling of the code. The Real-Time Systems Group of TU-Vienna participating in the Timing Analysis platform and the Programming Languages Group participating in the Compiler platform started a cooperation on developing such a flow information transformation framework. The research is funded by the Austrian Science Fund (Fonds zur Förderung der wissenschaftlichen Forschung) within the research project “Compiler-Support for Timing Analysis" (COSTA). COSTA started in July 2006 and has duration of three years. A first prototype that is performing source-to-source transformation of flow information has been already implemented.
http://ti.tuwien.ac.at/rts/research/projects/COSTA/

Common Flow Description Attributes (TU-Vienna/Real-Time Systems Group, TU-Vienna/Programming Languages Group, Mälardalen University, AbsInt, University of York)

To compute the WCET of a program in general, additional flow information is needed to guide the WCET computation. Due to different research approaches and framework architectures, there exist many different flow description languages for WCET analysis, making it hard to connect different components of WCET analysis frameworks together. Within ARTIST2 the partners TU-Vienna and Mälardalen University together with the industrial partners started an effort with the ambitious goal to define flow description attributes that serve as an exchange format among different tool components. As a result within year 3, existing flow description languages have been analysed and described by TU-Vienna. The next activity will be to set up a web page as an information exchange portal for a “challenge on flow description languages”. The research is partially funded at the TU Vienna by the Austrian Science Fund (Fonds zur Förderung der wissenschaftlichen Forschung) within the research project “Compiler-Support for Timing Analysis" (COSTA). COSTA started in July 2006 and has duration of three years. The next steps of the ARTIST2 partners working on the flow description attributes will be to collect feedback received from this challenge on flow description languages, and to define a common set of flow information. Based on this common set of flow information a set of properties to be embedded as attributes into the AIR format will be defined.
http://ti.tuwien.ac.at/rts/research/projects/COSTA/

Four partners of the team (Vienna, Mälardalen, Tidorum, and AbsInt) continued to work on path description attributes for AIR to arrive at a uniform notation. This work led to the conclusion that a more detailed study and comparison of the possible formats is necessary. This study and comparison was done by TU Vienna. The results of this study have been documented and published. As a next step, it is planned to set up a challenge of flow-description languages to get a broad feedback about proper solutions for that problem domain. (*Due to the need for further investigations, the integration of new path description attributes into AIR has been moved beyond the scope of ARTIST2*).

Evaluation of SWEET and aiT at Volvo CE (Mälardalen, Volvo CE and AbsInt)

This evaluation was performed by Dani Barkah, student at KTH (The Royal Institute of Technology, Stockholm), as a Master's Thesis work. Researchers from Mälardalen University supervised Barkah. The purpose of the work was to test the automatic flow analysis (as described in the ECRTS 2008 paper; see 2.4.2) of the prototype WCET tool SWEET (SWEdish Execution time Tool) from Mälardalen University on real industrial code. A second purpose was to compare the automatic flow analysis of SWEET to manual flow analysis made in an earlier Master's Thesis work (described in Daniel Sehlberg, et al. Static WCET Analysis of Real-Time Task-Oriented Code in Vehicle Control Systems. the ISoLA-06, Cyprus, Nov. 2006). Yet a purpose was to run the WCET tool aiT using the results from SWEET. The title of Barkah's Master's Thesis is "Evaluating program flow analysis for WCET calculations at Volvo CE". The work has been finished in August 2007.
Dani Barkah and Nils-Erik Bånkestad (head of the software development at Volvo CE in Eskilstuna, Sweden) visited AbsInt in April 2007 as a part of this work.
http://www.mrtc.mdh.se/index.php?choice=publications&id=1341

WCET Analysis and Certification of Automatically Generated Code (Mälardalen, CC-Systems AB and Tidorum)

This work was performed by the Mälardalen student Elie Assaf as a Master's Thesis work. The commercial static WCET Tool Bound-T was used to conduct a WCET analysis on parts of a C++ code that was automatically generated from the component-based tool Rhapsody. The C++ code was generated for a hard real-time application that runs on a bridge control panel installed in a Rolls Royce marine vessel. Researchers from Mälardalen University as well as Niklas Holsti from Tidorum AB, Finland, supervised Assaf. The title of Assaf's Master's Thesis is "WCET Analysis and Certification of Automatically Generated Code for CC-Systems AB". The work finished in August 2007.
http://www.mdh.se/ide/eng/msc/index.php?choice=show&id=0628

Conversion of the AbsInt AIR format to SWEET Format (Mälardalen and AbsInt)

This work is performed by the Mälardalen student Per Wolde as a Master's Thesis work. Researchers from Mälardalen University are supervising Per Wolde. This work will allow SWEET to analyse binaries for the NECV850 processor using AbsInt tools like aiT or "exec2crl". The chain of files and tools is: C-source → NEC_gnu_cross_compiler → exec2crl → crlreader → SWEET. Per Wolde (together with a group of MSc and PhD students from Mälardalen University) visited AbsInt in the winter of 2006/2007 as a part of this work. This is on-going work.
http://www.mdh.se/ide/eng/msc/index.php?choice=show&id=0536

Transformation of Flow Facts within Optimizations of a WCET-aware Compiler (Dortmund University)

Timing analysis relies on the presence of highly precise flow facts. Flow facts represent information about the possible flow of control through a program under analysis – e.g. iteration counts of loops. Usually, such information is provided by the designer who is fully responsible for its correctness. In the context of the WCET-aware compiler developed at Dortmund in the past years, flow facts can now be entered into the compiler within the source code to be processed by the compiler. These flow facts are analyzed and kept semantically correct during each transformation performed by the compiler. In particular, all flow facts are maintained and adjusted during all compiler optimizations, even if they heavily restructure the code. These automatically transformed flow facts are finally passed to the WCET analyzer aiT provided by AbsInt, in order to perform the actual WCET analysis. This achievement has taken away the burden from the designer to specify flow facts at the assembly code level. Now, the designer can annotate the source code, invoke the compiler, perform optimizations, and still obtains valid WCET results.
http://ls12-www.cs.uni-dortmund.de/

Compile-Time Decided Instruction Cache Locking Using Worst-Case Execution Paths (Dortmund University, AbsInt)

Caches are notorious for their unpredictability. It is difficult or even impossible to predict if a memory access results in a definite cache hit or miss. This unpredictability is highly undesired for real-time systems. The Worst-Case Execution Time (WCET) of software running on an embedded processor is one of the most important metrics during real-time system design. The WCET depends to a large extent on the total amount of time spent for memory accesses. In the presence of caches, WCET analysis must always assume a memory access to be a cache miss if it can not be guaranteed that it is a hit. Hence, WCETs for cached systems are imprecise due to the overestimation caused by the caches.
Modern caches can be controlled by software. The software can load parts of its code or of its data into the cache and lock the cache afterwards. Cache locking prevents the cache's contents from being flushed by deactivating the replacement. A locked cache is highly predictable and leads to very precise WCET estimates, because the uncertainty caused by the replacement strategy is eliminated completely.
In year 3 of ARTIST2, the lockdown of instruction caches at compile-time to minimize WCETs was explored. In contrast to the current state of the art in the area of cache locking, our techniques explicitly take the worst-case execution path into account during each step of the optimization procedure. This way, we can make sure that those parts of the code are locked in the I-cache that leads to the highest WCET reduction. The results demonstrate that WCET bound reductions from 54% up to 73% can be achieved with an acceptable amount of overhead required for the optimization and WCET analyses themselves.
http://ls12-www.cs.uni-dortmund.de/

Influence of Procedure Cloning on WCET Prediction (Dortmund University, AbsInt)

For the worst-case execution time analysis loops are an inherent source of unpredictability and loss of precision. This is caused by the difficulty to obtain safe and tight bounds on the number of iterations executed by a loop in the worst case. In particular, data-dependent loops whose iteration counts depend on function parameters are extremely difficult to analyze precisely. Procedure Cloning helps by making such data-dependent loops explicit within the source code, thus making them accessible for high-precision WCET analyses.
In year 3 of ARTIST2, we studied the influence of standard optimizations found in ordinary compilers on the program's WCET. We present the effect of Procedure Cloning applied at the source-code level on worst-case execution time. The optimization generates specialized

versions of functions being called with constant values as arguments. In standard literature, it is used to enable further optimizations like constant propagation within functions and to reduce calling overhead.
Our work shows that Procedure Cloning for WCET minimization leads to significant improvements. Reductions of the computed WCET bounds from 12% up to 95% were measured for real-life benchmarks. These results demonstrate that Procedure Cloning improves analyzability and predictability of real-time applications dramatically. In contrast, average-case performance as the criterion Procedure Cloning was developed for is reduced by only 3% at most. Our results also show that these WCET reductions only implied small overhead during WCET analysis.
http://ls12-www.cs.uni-dortmund.de/

Final Results

ARTIST Interchange Representation and Attribute Database: AIR (AbsInt)

Work on the AIR format continued. The format was extended and adapted to the needs of the partners. The attribute database was extended by new attributes as required.

Computation Semantics Representation: ALF (USaar, Tidorum, Mälardalen, AbsInt)

To support the flow analysis (loop bounds and infeasible paths) AIR must be able to represent the computations executed by the instructions in the program under analysis. CRL2, the basis of AIR, does not have an explicit, portable computation semantics representation. Work on such a representation started in Year 2 and continued in Year 3 and 4. It was decided that the ALF language from Mälardalen University meets the requirements, and should be used for this purpose. Mälardalen produced a specification of ALF, which finalizes the preliminary specification produced in Year 3.
http://www.mrtc.mdh.se/index.php?choice=publications&id=1351

Conversion of the AbsInt AIR format to SWEET Format (Mälardalen and AbsInt)

This work will allow SWEET to analyse binaries for the NECV850 processor using AbsInt tools like aiT or "exec2crl". The AIR format for NECV850 is translated into SWEET's internal code representation for low-level representation. The translator is close to final. AbsInt's existing StackAnalyzer for NECV850 was extended by a timing analysis component.

http://www.mdh.se/ide/eng/msc/index.php?choice=show&id=0536

Common Flow Description Attributes (TU-Vienna/Real-Time Systems Group, TU-Vienna/Programming Languages Group, Mälardalen University, AbsInt, University of York)

Defining common flow description attributes has shown to include more aspects than initially expected. Besides the flow description it is also necessary to describe other aspects like the memory layout, absolute timings, etc. We decided to postpone the final proposal of common flow description attributes to get more feedback from the community. To get more feedback, we proposed the WCET Annotation Language Challenge at the 7th International Workshop on Worst-Case Execution Time Analysis (WCET 2007) in Pisa, Italy. The website of the WCET Annotation Language Challenge has been set up at the address

http://costa.tuwien.ac.at/languages.html.

Meanwhile we published a paper that presents what we consider essential ingredients. This paper is intended to trigger concrete feedback. It is planned to post feedback from other people also on this web page to reflect the current status.

WCET Challenge 2008 (Tidorum, Mälardalen, USaar)

The structure of the Challenge, as a set of benchmark programs, analysis tasks, and results, was developed from the loose structure of the 2006 Challenge to a more detailed and precise form, shown on the Challenge Wiki site at http://www.mrtc.mdh.se/projects/WCC08/doku.php.

It was agreed that the Challenge should be considered a continuous process, allowing the addition of benchmark programs, analysis tasks, participants, and results at any time, but punctuated by an annual deadline at which a snapshot is taken and becomes the outcome of the Challenge for that year.

Most participants in the 2008 Challenge found a need for a standard, formal language to describe the execution flows assumed in the analysis tasks. This creates an interesting connection to the Annotation Language Challenge of the CoSTA project (http://costa.tuwien.ac.at/languages.html).

In addition to the ARTIST2 partners listed in the heading above, who were active in organizing the 2008 Challenge, several other ARTIST2 partners participated in the Challenge with their tools: TU Dortmund University, TU Vienna Real-Time Systems Group, TU Vienna Programming Languages Group, and University of York (through Rapita Systems).

Timing Analysis and Timing Predictability (USaar and AbsInt)

The notion of predictability of cache architectures has been clarified. This is the first precise notion of predictability found in the literature. It turns out that the cache-replacement strategy is the decisive characteristic for the predictability of architecture. 4 different replacement strategies were compared, and the LRU strategy was found to be optimal.

The PREDATOR project in the 7th Framework Programme attempts to reconcile performance and predictability. It has identified the PROMPT (PRedictability Of Multi-processor Timing) design rules for predictable multi-processor design. The first principles are to avoid interference on shared resources in the architecture and to allow the application designer the mapping of applications to target architecture without the introduction of new interferences that were not present in the application.

Parametric Timing Analysis (USaar and Mälardalen)

Timing analyses require that information such as bounds on the maximum numbers of loop iterations are known statically, i.e., during design time. Parametric timing analysis softens these requirements: it yields symbolic formulas instead of single numeric values representing the upper bound on the task's execution time. So, some input parameters to the program can remain unknown until the final use of the task. The developed analysis determines the parameters of the program, constructs parametric loop bounds, takes processor behaviour into account and attains a formula automatically.

Synergy between Code Synthesis and Timing Analysis (USaar and AbsInt)

One of the problems to be solved synergetically by code synthesis compiling, and timing-analysis is to support mode-specific timing analyses. Many embedded control systems have several operating modes with different timing requirements. Some operating modes are not explicitly specified on the model level or in comments on the C level. Saarland University in cooperation with Bosch (within the PREDATOR project), attempts to semi-automatically identify such operating modes. This would allow timing analysis to implement mode-specific execution-time bounds which can be used to improve schedulability of task sets.

WCET Analysis for Systems with Preemptive Scheduling (USaar and Absint)

Derivation of timing guarantees was extended in order to cope with the systems with preemptive scheduling. A new method to compute valid upper bounds on a task's worst case execution time (WCET) under preemption was proposed. This method approximates an optimal memory layout such that the set of possibly evicted cache-entries during preemption is minimized. This set then delivers information to bound the execution time of tasks under preemption in an adopted WCET analysis.

Transformation of Flow Information during Compilation (TU Vienna/Real-Time Systems Group + TUV/Programming Languages Group)

Flow information is used to guide the WCET analysis framework in finding the worst-case path. Such flow information may be calculated automatically or given as manual code annotations. Also the automatic calculation of flow information is easier at higher abstraction levels than the machine code. Thus, in both cases, the automatic calculation or the manual annotation a mechanism is required to transform the flow information down to the object code level where the WCET analysis is done. The initial work done by Kirner in his PhD thesis has been extended to the transformation of flow information at source-code level and interprocedural program optimization. The transformation of the flow information has been integrated into the interprecedural program transformation framework SATiRE, which connects other frameworks like ROSE or PAG. Adding the support of transforming flow information at the source-code-level has the advantage that most transformations that change the structure of the code can be already done in a rather compiler-independent way at the source code. Thus, on a compiler that does not support the transformation of flow information, one can deactivate all code optimisations In the compiler that change the structure of the control flow and use instead code transformations at the source-code level. Further information can be found at the web site of the CoSTA project: http://costa.tuwien.ac.at/languages.html.

WCET- and Memory Architecture-aware Compilation (TU Dortmund, AbsInt)

See the activity report for the Compilers Platform Activity.

WCET-aware Procedure Positioning and Cloning (TU Dortmund, AbsInt)

See the activity report for the Compilers Platform Activity.

Using Learning to Support the Development of Embedded Systems (York)

In 2007, the University of York was awarded 3.5 years of funding by EPSRC to investigate how novel techniques could be used as part of WCET analysis. The first part of the funding is for a post-doc to investigate how machine learning can be used to infer models of software and hardware as part of WCET analysis. To date two publications have been produced as part of this grant. The first paper (Bartlett 2008) builds on previous work that considers how Inductive Logic Programming can be used to establish relationships in loop guards to reduce the pessimism in program flow analysis. In particular the latest paper look at how the scalability of the work can be improved. The second paper, (Bate 2008), looks at how models of branch predictors can be determined by assessing program traces and then how this can be combined with static analysis. The second part of the funding is for a PhD student who will investigate which coverage metrics are most appropriate for WCET analysis and then how automated techniques can be used to generate the test cases needed.

Handling Timing Anomalies for Efficient WCET Analysis (USaar)

Abstractions employed for static timing analysis can lead to non-determinism that may require the analyzer to evaluate an exponential number of choices even for straight-line code. Pruning the search space is difficult because of the danger of "Timing Anomalies" where local worst-case choices may not lead to the global worst-case scenario. Earlier work explored the characterization and identification of such anomalies. Starting with the assumption that all reasonable abstractions of modern hardware do exhibit timing anomalies, we explored on the theoretical level, when and how the search space can be safely pruned to enable more efficient, yet safe, WCET analyses. To this end, we propose an approach that uses precomputed information to safely discard states in almost constant time.

The Impact of Timing Anomalies on WCET Analysis (TU Vienna)

There are two different phenomena that are called timing anomaly: the inversion of a local effect at the overall execution time and the amplification of a local effect at the overall execution time. We explored different the effects of timing anomalies at a theoretical level and analysed whether a stepwise timing analysis of hardware effects is still possible. We published first results where we described two different timing composition methods. The interesting result is that each timing composition method is safe if only one type of timing anomaly is present. However, in the case that both types of timing anomalies are present both techniques are not safe. These results provide more insights to the nature of timing anomalies and how they are challenging to WCET analysis.

Compilers (Platform)
Led by Sabine Glesner (TU Berlin)

Partner teams (leaders): Sabine Glesner – Technical University of Berlin (Germany), Rainer Leupers – RWTH Aachen University (Germany), Peter Marwedel – Dortmund University (Germany), Hans van Someren – ACE (The Netherlands), Reinhard Wilhelm – Saarland University (Germany).

Affiliated teams (leaders): Christian Ferdinand – AbsInt (Germany), Stylianos Mamagkakis, Prof. Francky Catthoor – IMEC vzw. (Belgium), Markus Schordan, Prof. Andreas Krall – TU Vienna (Austria).

Overview: Traditionally, timing analysis tools have been designed independently of compilers. It has now turned out that proceeding along this path would result in a duplication of efforts. Flow facts are available in compilers and need to be regenerated in timing analysis tools. Timing information is available in timing analysis tools and would be useful for timing-aware optimizations in compilers. Currently, compilers use very rough approximations of timing, if they use any timing model at all. As a result, the impact of certain transformations on run-time is frequently not known by compilers. Hence, the user has to follow a trial-and-error approach, experimenting with different compiler options and figuring out a suitable combination of them. However, even this time-consuming process cannot really minimize the execution time since options which might be good for some part of the code might lead to bad result for some other part of the code. A tight integration of timing models into compilers and their optimizations is urgently needed.

There is a general trend in the industry to replace non-programmable hardware accelerators (NPAs) with flexible reconfigurable cores, which have specialized resources and instructions dedicated to a class of applications. These reconfigurable cores create new challenges for embedded development tools and especially for compilers, and new challenge for processor architecture investigation tools.

Examples of configurable cores include Xtensa from Tensilica, ARC600 and 700 from ARC, CoreXtend from MIPS. Examples of flexible development tools are the Coware/LISATek processor and compiler designer based on CoSy Express, or the toolsets proposed by Tensilica and ARC for their core extension development.

Many applications in the embedded systems domain are both resource-restricted and safety-critical. This in turn requires compilers for embedded processors to be both efficient and correct. In cooperation between the Technical University of Berlin and ACE, verification methods and tools for compilers have been investigated.

Work in Year 1

Cooperation AbsInt – TU Vienna

Our goal for Year 1 was the integration of the Program Analysis Generator (PAG) of AbsInt in several platforms to share the same analysis in different infrastructures and leverage existing optimizations for evaluation. We created a tool, the PAG Interface Generator (PIG), to automate the PAG integration. The two different infrastructures which served as applications for the PAG integration by using PIG were ROSE and OCE/xDSPcore. ROSE is a source-to-source infrastructure that supports C++ (and Fortran in near future). The OCE/xDSPcore is the ATAIR open compiler infrastructure with a backend for digital signal processors.

Achievements for Year 1: Our goal for Y1 was the creation of a tool, PIG, to automate most aspects of the PAG integration and prove its usefulness by using it for integrating PAG in ROSE and OCE/xDSPcore. We have achieved both goals such that we can demonstrate the result by having a constant propagation analysis (as test) running in both environments.

Cooperation IMEC – University of Dortmund

The cooperation between IMEC and University of Dortmund resulted in the alignment of the research objectives for the steering of locality-improving loop transformations at the source code level. For this purpose, the control flow complexity of a given source code should be evaluated for steering the loop-transformations and evaluating their benefits and overheads, before actually compiling the resulting code on the target platform. The requirements (the WHAT specifications) to tackle this problem were defined. Based on the WHAT specifications, Dortmund looked at high-level control flow cost estimation approaches that could base its estimate when only the source code is available (without performing any compilation). At IMEC complementary actions had been started to see how this estimator can be integrated in a loop transformation framework project that had been started up earlier (prior to the start of ARTIST2) and that is now being extended for these high-level estimators.

Work in Year 2

-- Dortmund – AbsInt --

Design of a WCET-aware C Compiler

Based on the interface language CRL2 of AbsInt's timing analysis tool aiT, a successful integration of timing analysis into the compiler infrastructure of Dortmund University was achieved. This was done by automatically translating the assembly-like contents used in compilers to aiTs CRL2 format. Additionally, the results produced by the WCET analyzer aiT were automatically collected and re-imported into the compiler infrastructure. This way, precise timing information is available within a compiler for future optimization for the very first time. In

addition, a powerful mechanism was developed to attach not only WCET-related data to the compiler data structures, but also to store arbitrary information used by optimizations targeting different objectives than WCET. This approach will be useful in order to perform automated trade-offs between different optimization goals.

-- Source Code Transformation for WCET-Optimization --

The influence of the loop nest splitting source code optimization on the worst-case execution time (WCET) was examined. Loop nest splitting minimizes the number of executed if-statements in loop nests of embedded applications. It identifies iterations of a loop nest where all if-statements are satisfied and splits the loop nest such that if-statements are not executed at all for large parts of the loop's iteration space. Especially loops and if-statements of high-level languages are an inherent source of unpredictability and loss of precision for WCET analysis. As a consequence, the optimization achieves a significantly more homogeneous control flow structure. Additionally, the precision of the optimization algorithms led to the generation of very accurate high-level flow facts. All together, considerable reductions of WCET were achieved by the source code optimization.
http://ls12-www.cs.uni-dortmund.de/research/C2C

-- ACE – Aachen --

-- Optimization of Conditional Execution in CoSy --

A dynamic programming algorithm is being implemented and tested on a number of different architectures to validate its behavior with real world code and current high-end industrial processors.

A prototype comprising a set of optimization engines and compilers has been constructed.

No-one has successfully been able to find a formalism or generate tools which facilitate generic retargeting of these algorithms.

-- TU Berlin – ACE --

TU Berlin, who joined the ARTIST2 NoE as an affiliated partner in Y2 and became a core partner in Y3, has worked on the verification as well as on the development of optimizing compiler transformations and machine code generation. Especially in safety-critical applications in the embedded domain, compiler transformations must be both optimizing and correct. Hence, verification is necessary to ensure that transformations indeed preserve program semantics during compilation. Within ARTIST2, the focus is on the development of automated checkers that, for a particular compiler run with its source and target program, make sure that both programs are indeed semantically equivalent. As a starting point, the verification and development of checkers for loop transformations based on unimodular transformations has been investigated.

-- Dortmund – IMEC --

The main technical outcome of the Dortmund-IMEC collaboration has been an agreement on the basic guidelines for the source to source transformations regarding static and dynamic optimizations (at design time and at run time respectively). These optimizations will target the loop transformations and memory assignment of statically and dynamically allocated data in complex memory hierarchies. The collaboration is mainly based on synchronized, individual work of each of the two partners and aims on common work through PhD research.

-- AbsInt – TU Vienna --

Extension of the ROSE-PAG integration from C to C++ and Implementation of Alias Analysis.

The ROSE-PAG integration achieved in Y1 for C was substantially extended to cover full C++ (only excluding Exceptions). This includes handling of templates, virtual methods, short-circuit evaluation in conditions, resolving overloaded functions, C++ name spaces, constructor and destructor calls. An intra-procedural shape analysis, published by our cluster partner Reinhard Wilhelm, was implemented using PAG. We extended the analysis to an inter-procedural shape analysis. The results of the analysis can be written to an external file and visualized using the tool AiSee.

Infrastructure for high-level specification of C++ program analyses

With the integration of PAG in ROSE, an infrastructure is available that permits using a high-level language for specifying an abstract interpretation of C++ programs. ROSE uses the EDG front end for parsing C++ and offers a powerful interface for accessing and transforming the abstract syntax tree (AST). The decorated AST offers the full type information of C++ input program and the PAG-ROSE integration permits using this type information in the PAG specification (e.g. for virtual method resolution).

Difficulty: Handling of the wide range of programming constructs of a general-purpose language

C++ has such a rich set of programming constructs that research prototypes of analyses often only consider subsets of C/C++. Our goal was to create an infrastructure that permits performing research on real-world programs. In particular, the interface between PAG and ROSE required a careful design, such that we can maintain updates of ROSE and PAG, but keep the required changes in existing analysis specifications at a minimum. Our approach is grammar based and permits the generation of the used design patterns, glue code (between ROSE and PAG), and implementations of the required interfaces.

Work in Year 3

SIMD & Conditional Execution Support in CoSy (Aachen, ACE)

- **SIMD retargeting**: A retargeting formalism for the SIMD optimization developed earlier has been devised. The SIMD instructions can now be described within CoSy's natural code generator description format. To achieve this, adaptations to the backend generator were necessary and improvements in data dependency analysis were devised.
- **SIMD enabling loop transformations**: Several loop transformations, such as strip mining and loop peeling, that increase the number of possible SIMD operations have been implemented in the CoSy system. They interact with the SIMD optimizer and consult the code generator description in order to assess the benefit of each transformation.
- **Conditional execution retargeting**: The retargetability of the conditional execution optimization previously done has been shown by adding support for a commercial, CoSy based, compiler. The results of this work by a student from Aachen led to several improvements in the cost calculation and to beneficiary transformations.

Transformation of Flow Facts within Optimizations of a WCET-aware Compiler (Dortmund University)

Timing analysis relies on the presence of highly precise flow facts. Flow facts represent information about the possible flow of control through a program under analysis – e.g. iteration counts of loops. Usually, such information is provided by the designer who is fully responsible for its correctness. In the context of the WCET-aware compiler developed at Dortmund in the past years, flow facts can now be entered into the compiler within the source code to be processed by the compiler. These flow facts are analyzed and kept semantically correct during each transformation performed by the compiler. In particular, all flow facts are maintained and adjusted during all compiler optimizations, even if they heavily restructure the code. These automatically transformed flow facts are finally passed to the WCET analyzer aiT provided by AbsInt, in order to perform the actual WCET analysis. This achievement has taken away the burden from the designer to specify flow facts at the assembly code level. Instead, the designer now can annotate the source code, invoke the compiler, perform optimizations, and still obtains valid WCET results. http://ls12-www.cs.uni-dortmund.de/

Compile-Time Decided Instruction Cache Locking Using Worst-Case Execution Paths (Dortmund University, AbsInt)

Caches are notorious for their unpredictability. It is difficult or even impossible to predict if a memory access results in a definite cache hit or miss. This unpredictability is highly undesired for real-time systems. The Worst-Case Execution Time (WCET) of software running on an embedded processor is one of the most important metrics during real-time system design. The WCET depends to a large extent on the total amount of time spent for memory accesses. In the presence of caches, WCET analysis must always assume a memory access to be a cache miss if it cannot be guaranteed that it is a hit. Hence, WCETs for cached systems are imprecise due to the overestimation caused by the caches.

Modern caches can be controlled by software. The software can load parts of its code or of its data into the cache and lock the cache afterwards. Cache locking prevents the cache's contents from being flushed by deactivating the replacement. A locked cache is highly predictable and leads to very precise WCET estimates because the uncertainty caused by the replacement strategy is eliminated completely.

In year 3 of Artist2, the lockdown of instruction caches at compile-time to minimize WCETs was explored. In contrast to the current state of the art in the area of cache locking, our techniques explicitly take the worst-case execution path into account during each step of the optimization procedure. This way, we can make sure that always those parts of the code are locked in the I-cache that lead to the highest WCET reduction. The results demonstrate that WCET reductions from 54% up to 73% can be achieved with an acceptable amount of CPU seconds required for the optimization and WCET analyses themselves. http://ls12-www.cs.uni-dortmund.de/

Influence of Procedure Cloning on WCET Prediction (Dortmund University, AbsInt)

For the worst-case execution time analysis, especially loops are an inherent source of unpredictability and loss of precision. This is caused by the difficulty to obtain safe and tight information on the number of iterations executed by a loop in the worst case. In particular, data-dependent loops whose iteration counts depend on function parameters are extremely difficult to analyze precisely. Procedure Cloning helps by making such data-dependent loops explicit within the source code, thus making them accessible for high-precision WCET analyses.

In year 3 of Artist2, we studied the influence of standard optimizations found in ordinary compilers on the program's WCET. We present the effect of Procedure Cloning applied at the source-code level on worst-case execution time. The optimization generates specialized versions of functions being called with constant values as arguments. In standard literature, it is used to enable further optimizations like constant propagation within functions and to reduce calling overhead.

Our work shows that Procedure Cloning for WCET minimization leads to significant improvements. Reductions of the WCET from 12% up to 95% were measured for real-life benchmarks. These results demonstrate that Procedure Cloning improves analyzability and predictability of real-time applications dramatically. In contrast, average-case performance as the criterion Procedure Cloning was developed for is reduced by only 3% at most. Our results also show that these WCET reductions only implied small overhead during WCET analysis. http://ls12-www.cs.uni-dortmund.de/

Optimization and Verification in Compilers (TU Berlin, ACE)

TU Berlin works on the verification as well as on the development of optimizing compiler transformations and machine code generation. Especially in safety-critical applications in the embedded domain, compiler transformations must be both optimizing and correct. Hence, verification is necessary to ensure that transformations indeed preserve program semantics during compilation. Within ARTIST2, the focus is on the development of automated checkers that, for a particular compiler run with its source and target program, make sure that both programs are indeed semantically equivalent

TU Berlin has established a platform for a research compiler, using the compiler tool CoSy provided by ACE. We developed a backend specification for the Intel Itanium processor, which gave us an industrial-strength compiler for this architecture (the current SPEC benchmark suite CPU2006 is compiled). The challenges in establishing this platform have been to exploit the special features of the Itanium, e.g. predication. On top of that, first optimizations have been developed. These optimizations mainly aim at reducing the impact of the memory wall on program performance, which is drastic on modern VLIW processors like the Itanium.

Furthermore, we investigate how verification techniques can be brought into practice, in order to ensure that the code generated by the compiler is correct. We considered the scheduling phase of a compiler, which rearranges the instructions of a program in order to improve runtime performance. We formalized the scheduling phase in the theorem prover Isabelle/HOL and derived a criterion that is necessary for correctness of the scheduled code. From this criterion, we automatically generated the core of a checker (facilitated by the code generation capabilities of Isabelle) which can be used to augment any existing scheduler. By this, we discovered a bug in the scheduler of the popular GNU assembler, which lead to possibly incorrect programs. http://www.pes.cs.tu-berlin.de/

Optimized Dynamic Memory Allocation (IMEC vzw.)

The main technical achievement of IMEC (in the context of the collaboration with Dortmund Uni.) is the realization of multiple fine-grain dynamic memory allocation design options in software modules of the standard compiler library (e.g., *libc*), which can be parameterized and combined in many ways. Then, we extract application specific information (i.e., Software Metadata) and memory hierarchy specific information (i.e., Hardware Metadata) from the embedded system design. All this metadata information is exploited by our tools, which parameterize and combine the aforementioned dynamic memory allocation modules using energy efficiency criteria. The final result is the construction of a unique dynamic memory allocator, which is compiled with the software application and utilizes a unique combination design options. This unique design fine tunes its energy consumption management according

to the unique characteristics of the software application and the memory hierarchy of the embedded system.

Automatic Source-Code Annotation (TU Vienna, AbsInt)

Source-Code annotations allow for providing additional semantic information. Our goal was to support programmers in providing the information that the analyzer can determine automatically in a readable form, such that the need for user-defined annotations is minimized. We have created a general mechanism that allows for annotating source codes automatically with a textual annotation representing the results of arbitrary analysis results. The annotation mechanism is implemented as a source-to-source transformation at the statement level, generating analysis information either as comments or pragmas at the corresponding statement positions in the C/C++ source-code. A first application of this mechanism allows us to generate may-alias and must-alias annotations based on the results of a shape-analysis.

External Program Representation for Tool Interoperability (TU Vienna, AbsInt)

An external program representation permits to build tool chains for program analysis and transformation. We have designed and implemented an external representation for C programs. The connection has been established in both directions, for generating an external C representation as well as reading in the external representation. For the external format we have chosen Prolog syntax, because it is suitable for querying programs and specifying transformations at a high-level. This new feature has already been used in cooperation with the CoSTA project for specifying the transformation of WCET annotations according to loop optimizations performed with our LLNL-ROSE loop optimizer at the C code level.

The tools LLNL-ROSE, the ROSE loop optimizer, the Program Analysis Generator (PAG), and the generator and parser for the external program representation have been integrated within the Static Analysis Tool Integration Engine (SATIrE), which aims at offering the combined features of all tools to the user through a single interface and a set of new specialized tools. The full C++ language has been addressed, in particular virtual methods, templates, constructor/destructor calls, function pointers, etc. – only exceptions are not addressed yet. ROSE permits generating C++ code and lowered C code. The generated code can serve as input to ACE's compiler for generating optimized machine code. http://www.complang.tuwien.ac.at/markus/satire/

Final Results

Design of a Static Loop Analyzer (TU Dortmund)

Knowledge about the loop iteration counts is mandatory for a large number of different program analyses (e.g. loop transformations like loop unrolling or loop tiling). Furthermore, it is impossible to derive WCET information statically without knowing iteration counts of loops. In the past, this information was collected and annotated manually which was a very time consuming and error-prone job. By integrating a loop analyzer into the existing WCET-aware compiler framework developed at TU Dortmund, it is now possible to produce WCET-optimized code fully automatically.

Static loop bound analysis has the same complexity as the halting problem. Because of this, it is impossible to get exact results in an acceptable amount of time for general problems. The only way to do loop analysis is to reduce complexity by applying approximations. This is the reason why our loop analyzer is based on an enhanced version of the classical paradigm of Abstract Interpretation which is a mathematical framework to obtain sound approximations of the original problem.

The main disadvantage of conventional Abstract Interpretation is that it is still a very time consuming analysis technique if the results should not be too pessimistic. It analyses loops by evaluating every iteration on its own. By integrating a new static polytope-based loop evaluation inside the original Abstract Interpretation, it was possible to reduce the required analysis time without getting inferior results. If a loop is reached during fixed-point iteration of Abstract Interpretation, it is checked whether the new polytope-based loop analysis is applicable. If so, the loop body is analyzed only once and the iterative calculation of Abstract Interpretation can be skipped. If polyhedral analysis is not applicable, the loop is analyzed conventionally.

To analyze loops with the new polytope-based technique, several constraints concerning the loop's header and body have to be met. These constraints require that the body of the loop is not too complex. This is achieved by integrating another technique called Program Slicing. Program Slicing determines all expressions that are superfluous for the current analysis so that their evaluation can be skipped.

To demonstrate the applicability of the developed loop analyzer, we tested it using 96 benchmarks containing more than 700 loops. These benchmarks were taken from the benchmark suites MRTC, DSPStone, MiBench, UTDSP and MediaBench. Related to these benchmarks, 99.15% of all loops could be analyzed successfully. Furthermore, 95.62% of all loops could be analyzed without any overapproximation thus leading to exact results. Concerning the running time of our analysis, we could accelerate the time consumed by conventional Abstract Interpretation up to a factor of 100 for some benchmarks it the combination of Program Slicing and polyhedral loop analysis was enabled.

Another fact showing that our loop analyzer is also applicable to real-world problems is that it was the only tool which was able to answer all questions related to flow fact during the WCET tool challenge 2008, organized by the Timing Analysis cluster of Artist2.
http://ls12-www.cs.tu-dortmund.de

WCET-aware Procedure Positioning and Cloning (TU Dortmund, AbsInt)

In year 3 of Artist2, we studied the influence of the standard compiler optimization Procedure Cloning on the worst-case execution time and could show that WCET reductions of up to 95% could be achieved. The main drawback of this optimization is its heavy code size increase. In year 4 we continued our work on Procedure Cloning and extended the optimization by WCET concepts. In a first extension, the novel WCET-driven Cloning focuses on the optimization of those functions that promise the highest WCET improvement. We achieve WCET reductions of 64.2%, while restricting the average code size increase to 22.6% which is a significant improvement compared to the average code size increase of more than 80% in the previous work. In a further extension, we combined our WCET-driven Procedure Cloning with a smart positioning of the newly created function clones. By placing the clones and their callers contiguously in memory, cache conflict misses can be reduced. Results show that the WCET can be reduced by 7% on average when Cloning is combined with a sophisticated positioning of the cloned functions.

Moreover, the compiler optimization Procedure Positioning was exploited for a WCET reduction. Procedure Positioning is a well known optimization aiming at the improvement of the instruction cache behavior. In the past, this technique was based on profiling information to improve the average-case performance of the program. We exploited the ideas behind Procedure Positioning for an effective WCET minimization. Our optimization operates on a call graph which is based on WCET information to find the function candidates for a contiguous memory allocation that promise the highest WCET reduction. Thus, our WCET-centric call graph is more reliable than previous approaches since it is valid for all program and all input data. We developed two types of the WCET-aware Positioning, an effective greedy and a fast

heuristic approach. Results on real-world benchmarks show that WCET reductions of 10% on average could be achieved while the average-case execution time (ACET) was decreased by 2% on average. This emphasizes the importance of compiler optimizations tailored towards an effective WCET reduction which can not be accomplished with standard ACET optimizations that take decisions based on other cost models. http://ls12-www.cs.tu-dortmund.de

WCET- and Memory Architecture-aware Compilation (TU Dortmund, AbsInt)

In year 4 of Artist2, we studied the influence of scratchpad memory allocation techniques on worst-case execution times. Here, we developed integer linear programming models in order to decide which parts of a program's code or data should be moved onto the highly predictable scratchpad memory. In contrast to the current state of the art in this area, we developed integrated models explicitly taking the critical path of a program defining its WCET into account. The optimizations developed during year 4 of Artist2 pre-compute the scratchpad contents during compilation time. During runtime of the optimized programs, the scratchpad contents remain unchanged. The results demonstrate that average WCET reductions of 11.6% can be achieved by simply moving parts of the global data of 30 representative benchmarks onto the scratchpad. First experiments with scratchpad memory allocation of program code show that WCET reductions of more than 50% can be achieved for several benchmarks. Unlike previously published approaches, our ILP formulations scale well so that our techniques only require a few CPU seconds to solve the optimization problem.

In addition to scratchpad memory allocation, we investigated WCET-aware register allocation. Within the compiler community, register allocation is considered the most important optimization since it leads to an optimized use of processor registers which are by far the most efficient memories of an entire system. Traditionally, register allocation relies on graph coloring approaches which apply some simple heuristics to decide where so-called spill code (i.e. load/store instructions swapping registers in and out to main memory) has to be inserted. Since these traditional techniques are timing-unaware, they may lead to spill code generation along a program's critical path defining the WCET. By making a graph coloring register allocator WCET-aware, large average WCET reductions of 27.3% were measured for a total of 28 representative benchmarks. It is worthwhile mentioning that still huge gains can be achieved even in such well-established areas like register allocation. http://ls12-www.cs.tu-dortmund.de

Optimization and Verification in Compilers (TU Berlin, ACE)

TU Berlin works on the verification as well as on the development of optimizing compiler transformations and machine code generation. Within ARTIST2, the focus is on the development of automated checkers that, for a particular compiler run with its source and target program, make sure that both programs are indeed semantically equivalent

TU Berlin has established a platform for a research compiler, using the compiler tool CoSy provided by ACE. We developed a backend specification for the Intel Itanium processor, which gave us an industrial-strength compiler for this architecture (the current SPEC benchmark suite CPU2006 can be compiled). We investigated the development of novel compiler optimizations to mitigate the impact of the memory wall on program performance, which is drastic on modern VLIW processors like the Itanium. One result lies in the development of speculative compiler optimization (Speculative register promotion for global variables), which reduces the number of loads induced by the use of global variables and thereby leads to a better memory performance. We could show that this optimization leads to notable performance improvement for the SPEC CPU2006 benchmark suite. Second, we also considered how memory accesses in general can be optimized. To this end, we developed a speculative optimization, which targets all kinds of memory accesses. It requires precise information about the dependencies amongst memory accesses, to decide whether or not a given optimization is beneficial. However, state-of-the-art alias analyses are too imprecise for that aim. As a consequence, we

propose to use statistical machine learning techniques to yield predictors, which can act as heuristics to determine the dependency amongst memory accesses. In our experiments, we could show that with the trained predictors, the dependencies could be efficiently and precisely predicted for unseen programs.

Furthermore, we investigated how the code generator phase, which is a crucial compiler optimization, can be verified formally. We considered code generators that are based on bottom-up rewriting, which is the most common technique in current compilers. The machine code is generated from the intermediate representation by applying a set of rules that specify how a construct on the abstract level can be mapped to the machine level. The correctness of these rules is necessary for the code generator to be correct. We developed a formal model for a subset of the compiler intermediate representation from the CoSy-Compiler and for a subset of the Itanium assembler. On top of this, we specified code generator rules and proved them correct. Besides the correctness proofs, we developed proof engineering techniques that help to handle the size of formalizations for complete rule sets. We investigated which strategies known from software engineering can be applied in the area of formal verification, too. In addition, we investigated methods and have preliminary results concerning proof automation. We could show correctness for a subset of the rules used in the specification for the Itanium. Besides, our results make the verification of complete code generator specifications more realistic. http://www.pes.cs.tu-berlin.de/

Source-To-Source WCET Analysis (TU Vienna, AbsInt)

TU Vienna and AbsInt work on scalable source-level analysis and annotation-based timing analysis methods. The presentation of analysis results and the iterative enhancement by the user with expert knowledge about the timing behavior of a given system is incorporated by allowing round-trip engineering of timing analyses. The infrastructure SATIrE allows building analyzers that take source-code annotations as additional input as well as automatically generate output as annotations. The iterative application of this approach increases productivity, by requiring the user only to annotate the timing relevant information that is not automatically computed. The integration of PAG was fundamental in investigating scalability and precision of analyses. The automatic annotation of programs with PAG analysis results is possible because of its combination with the LLNL-ROSE C/C++ backend in SATIrE. The importance of user-readable analysis results as annotations has also fostered the use of PAG and SATIrE in teaching program analysis at several universities in Y3 and Y4.

For timing analysis various supporting analyses are necessary. In Y4 SATIrE was enhanced with a Steensgaard-style points-to analysis. This analysis partitions a program's objects (variables and dynamically allocated memory regions) into equivalence classes, and models which classes may contain pointers to which other classes. Members of structures are treated as individual objects unless accesses through pointers of incompatible type make it necessary to collapse structure fields. The analysis runs in almost linear time in the size of the program, allowing it to scale to very large input programs.

Further more, TuBound, was created with SATIrE. It performs an inter-procedural context-sensitive interval analysis with PAG and computes loop bounds for specific loop patterns in Prolog. For loops where a loop bound cannot be established, annotations can be provided by the user. The implemented algorithm for loop bounds was evaluated with the Mälerdalen Benchmark suite. TuBound also participated in this year's WCET Tool Challenge 2008.
http://www.complang.tuwien.ac.at/markus/satire/

Retargetable Code Optimizations (RWTH Aachen, ACE)

The cooperation between ACE and Aachen on retargetable code optimizations has been continued. The major problems left were complete integration of the conditional execution

prototype into CoSy and further development on the analysis framework necessary to facilitate compilation for SIMD architectures.

The conditional execution engines have been extended by a strong retargeting formalism. Several compiler passes and a few extensions to the backend description used in CoSy have been devised to achieve this goal. This work has by now been productized in the CoSy release.

Work continues in two directions. Improved loop analysis infrastructure is developed by at ACE by a student from Aachen. The goal of this project is to deliver the information necessary for advanced optimizations like vectorization. Also a cooperation to make efficient code generation for clustered VLIW processors available is being planned.

ARTIST Interchange Representation and Attribute Database: AIR (AbsInt)

Work on the AIR format continued. The format was extended and adapted to the needs of the partners. The attribute database was extended by new attributes as required.
http://www.theiling.de/absint/attrdb.fcgi

Analyses on C source code (TU Vienna, AbsInt, Usaar)

AbsInt and TU Wien continued working on the Static Analysis Tool Integration Engine, SATIrE, which connects LLNL-Rose with AbsInt's Program Analysis Generator (PAG). In year four, SATIrE was improved and made operational. In cooperation with Saarland University, it was successfully applied to the analysis of loop bounds and function pointers.
http://www.complang.tuwien.ac.at/markus/satire/

2.2.4 Execution Platforms Cluster

This cluster is composed of the following activities:

System Modelling Infrastructure (Platform)
Led by Jan Madsen (Denmark Technical University)

Partner teams (leaders): Petru Eles – ESLAB, Linköping University (Sweden), Rolf Ernst – IDA, TU Braunschweig (Germany), Luca Benini – Micrel Lab, University of Bologna (Italy), Jan Madsen – IMM, Technical University of Denmark (Denmark).

Affiliated teams (leaders): Roberto Zafalon – STM (Italy, Salvatore Carta - University of Cagliari (Spain), Magnus Hellring – Volvo (Sweden), Bjørn Sand Jensen – Bang & Olufsen ICEpower (Denmark), Rune Domsteen – Prevas (Denmark).

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

The main objectives are:

- Integrate ongoing research efforts on infrastructure modelling
- Replacing prototyping of hardware
- Reducing the cost and time required for designing embedded systems
- Tackling the growing complexity of embedded systems

Work in Year 1

-- Simulation-based modeling (month 0 – 6) --

MPARM (UoB)

UoB has devoted quite a lot of effort to developing a simulation environment for multi-processor systems on chip, called MPARM, and augmenting its modelling capabilities so to be able to simulate real-life systems. The work initially addressed the integration of IP cores into the simulation environment, and therefore required to cope with issues such as IP core wrapping and standard interfacing, synchronization with the system simulation engine, development of non-functional models (e.g., power consumption) and integration of such models in the functional simulation. Moreover, inter-processor communication and synchronization mechanisms have been developed, modelled and explored via functional simulation. The work on IP core integration, communication and synchronization has led to a multi-processor system-on-chip simulation environment with unprecedented modelling capabilities, which can be effectively deployed for design space exploration at a high level of accuracy. In fact, the cycle-accuracy of all component models has been preserved throughout the development and integration process.

MPARM existed prior to ARTIST2, but it has been extended and refined during the first year of ARTIST2. The work contributed by ARTIST2 is the development of the entire traffic tracing and the transaction extraction facilities, which have been instrumental to integrate the traffic generators (see next paragraph) with MPARM. On the modelling side, ARTIST2 has

contributed with the development of a few communicaiton intensive benchmarks.for comparing different NoC technologies and strategies.

Traffic generators (UoB, DTU)

Traditionally, synthetic traffic generators have been used to overcome the more realistic development scenarios in the industry, where the parallel development of components may cause IP core models to be still unavailable when tuning the communication architecture. However, target applications increasingly present non-trivial execution flows and synchronization patterns, especially in presence of underlying operating systems and when exploiting interrupt facilities. This property makes it very difficult to generate realistic test traffic. For this reason, UoB and DTU have established a cooperation to realistically render SoC traffic patterns with interrupt awareness. The proposed methodology was extensively validated by showing cycle-accurate reproduction of previously traced application flows.

The entire work on developing, integrating and testing the traffic generators are contributed by ARTIST2. This has been done in a close cooperation between UoB and DTU.

ARTS (DTU)
ARTS is a system-level heterogeneous multiprocessor System-on-Chip (MPSoC) modeling framework which allows for designer-driven design-space exploration of heterogeneous MPSoC platform architectures and co-exploration of cross-layer dependencies. In particular, the consequences of different mappings of tasks to processors – software or hardware, the effects of different RTOS selections – scheduling, synchronization and resource allocation policies, and the effects of different network topologies and communication protocols. As both ARTS and MPARM are based on SystemC and interface components using OCP, initial tempts to interface the two models have been taken.
ARTS existed prior to ARTIST2, but it has been extended and refined during the first year of ARTIST2. The work contributed by ARTIS2 are the development of an OCP-based interface between computaion and communication components. This has required a major rewrite of the modelling core of ARTS as well as an extension of the basic ARTS model to include IO tasks, modelling device drivers, and IO devices, modelling the interface devices.

-- Formal modeling (month 0 – 6) --

Response time analysis (TU Braunschweig)

The aim has been to consider remote memory accesses in worst-case response times. In our extended task model, multiple remote transactions during the task's execution are taken into account. This extension allows the straightforward modelling of many applications and increases the applicability of formal analysis methods to many real-world architectures containing shared memory.

The task model extension that allows to include the effect of remote memory accesses on worst-case response time analysis is entirely contributed by ARTIST2.

Sensitivity analysis (TU Braunschweig)

Sensitivity analysis and flexibility optimization is used to reduce the design risk of critical components and increase design robustness. Sensitivity analysis allows the system designer to evaluate the flexibility of a given system, and thus to quickly assess the system-level impact of changes in performance properties of individual hardware and software components. If for example, the integration of a supplied IP component into the system in its original configuration

results in a non-working system, the designer can easily determine if the reconfiguration of the system is possible so that all system constraints are satisfied.

The Artist2 contribution to sensitivity analysis is the formal framework to define sensitivity and robustness measures which can then be used in design space exploration and robustness optimization.

-- Simulation-based modeling (month 6 - 12) --

MPARM (UoB)

The system modelling effort in MPARM has not only been concerned with the hardware architecture, but also the software infrastructure. A middleware layer has been designed with the objective of abstracting software developers from low level implementation details such as memory map, address of memory mapped slave devices, management of synchronization and inter-processor communication mechanisms, shared memory allocation and de-allocation, etc. This work has been performed with the cooperation of associated member IMEC.

The work on the middleware layer performed in cooperation with IMEC is entirely contributed by ARTIST2.

Trafiffic generators (UoB, DTU)

The cooperation between UoB and DTU to establish a reactive traffic generator model has been continued in this period. The generator model has been extended to deal with more complex and more realistic events, such as OS-driven interrupt handling mechanism, and therefore mimic non-trivial execution flows and synchronization patterns.

The work on extending the traffic generators to handle interrupts and synchronization is entirely contributed by ARTIST2.

Distributed embedded systems for automotive applications (Linköping University, DTU)

The work of Linkoeping Univeristy aims at implementing a simulator for distributed embedded systems for automotive applications. The starting point for this work is the multiprocessor simulation environment (ARTS) developed at DTU. The work is performed in cooperation between the DTU and Linkoping groups. A student from DTU has made a short visit to Linkoping in June 2005.

The entire work on developing a simulator for distributed embedded systems for automotive applications based on ARTS is contributed by ARTIST2.

ARTS (DTU)

The work on the ARTS modeling framework has been extended to include more details of the platform. This has been done in order to target two different types of platforms: MPSoC particularly for multimedia applications, and wireless sensor networks. On the application side, DTU and Linköping University have started a cooperation aimed at extending ARTS to be able to simulate distributed embedded systems for automotive applications.

The cooperation between DTU and Linkoping is done entirely within ARTIST2. The extensions related to multimedia applications, including the setup and experimentation wich a wireless multimedia terminal is contributed ARTIST2. The work done on extending the ARTS modeling framework to support modelling of wireless sensor networks has been done within the Hogthrob project, which is one of the main sources of fundings for the system modelling infrastructure activity.

-- Formal modeling (month 6 – 12) --

Modelling power consumption with SymTA/S (TU Braunschweig)

Recently the SymTA/S tool was extended in cooperation with Prof. Sharon Hu, University of Notre Dame, USA, to model and analyse the power consumption of complex heterogeneous embedded systems. Power aspects represent, besides performance issues, a critical problem during implementation and integration of complex systems. Currently we are working on power optimization techniques (heuristic and stochastic) based on the developed power models.

SymTA/S existed prior to ARTIST2. The work on extending the basic model of SymTA/S to handle power consumption which has been done in a cooperation between TU Braunschweig and University of Notre Dame is contributed by ARTIST2.

Integration (TU Braunschweig, ETHZ)

Initial attempts to integrate the event-stream formalism from TU Braunschweig and the real-time calculus from ETHZ have been established. The real-time calculus has been embedded into the SymTA/S front-end, allowing designers to expres both models within the same environment.

The cooperation between TU Braunschweig and ETHZ on embedding the real-time calculus within the SymTA/S front-end is entirely contributed by ARTIST2.

Work in Year 2

Simulation platform for distributed embedded systems (University Linköping)

A simulation environment is designed and implemented for distributed real-time systems such as those used in automotive applications. The ARTS environment, developed at DTU and targeting System on chip applications, has been used as a starting point by the Linköping team.

In Year2 the following work has been done:

- The implementation of the environment has been finalized and new protocols, such as Flexrey have been implemented;
- Theoretical investigation regarding anomalies and sensitivity in distributed real-time systems has been performed, with results that will help to improve the efficiency of the simulator in detecting close to worst case behavior. This is important when using the simulator for evaluation of the pessimism of certain schedulability analysis approaches. This work is done in interaction with the Braunschweig group.
- Implementation of real-life applications from our industrial partners at Volvo.

First publication is planned for the next year.

Modeling and response-time/buffer analysis for NoC (University Linköping)

The Linköping group has developed a system model, based on which worst case response times and worst case buffer need for hard real-time applications implemented on NoCs can be calculated. On top of this analysis approach, an optimization tool for buffer space minimization has been implemented, for real-time NoC applications.

Modelling and formal timing analysis of shared memory accesses (TU Braunschweig)

We have continued to investigate design paradigms of MPSoC architectures. As opposed to distributed systems, a common feature here is the use of a shared memory that is accessed from each processor, introducing conflicts on the memory and interconnects. System designers often implement latency-hiding techniques to reduce the effect of waiting for data, by allowing frequent context switches to tasks that are ready.

We have systematically identified dependencies in such systems that have an influence on design properties such as end-to-end delays. Using this in [SIE06], we were able to show that the technique for latency hiding can bear unwanted results for critical worst-case response time scenarios.

We have further investigated formally the timing of multiple coinciding memory accesses. Previous approaches had to assume a worst-case timing for each individual memory access. Due to large timing variations, this leads to a large deviation of analysis result and actual behaviour. In [SISE06], we presented a new way to express and calculate total latency of multiple events with much higher accuracy, leading to improve worst-case response time estimates.

Integration of formal SDF analysis techniques into the SymTA/S framework (TU Braunschweig)

Standard event models represent key integration aspects and hide complexity of local scheduling analysis algorithms. Thus, they are a suitable abstraction to integrate different models of computation into the SymTA/S framework. Recent work at IDA has produced a methodology to embed the analysis of SDF Graphs [Lee/Messerschmitt] into the SymTA/S framework (paper submitted for review at DATE07).

Integrating SDF models into the SymTA/S framework required corner-case evaluation of SDF graphs to construct event models describing their timing behaviour. Also, notions for path related metrics like latencies were defined and algorithms for computing their upper and lower bounds were proposed.

SDF Graphs are especially suited for describing data transforming applications like filters. Integrating their analysis into the SymTA/S framework significantly enlarges its application domain and improves the analysis results i.e. in the field of filter applications.

Multi-dimensional sensitivity analysis (TU Braunschweig)

The robustness of an architecture to changes is a major concern in embedded system design. Robustness is important in early design stages to identify if and in how far a system can accommodate later changes or updates or whether it can be reused in a next generation product. Robustness can be expressed as a "performance reserve", the slack in performance before a system fails to meet timing requirements. This is measured as design sensitivity.

Due to complex component interactions, resource sharing and functional dependencies, one-dimensional sensitivity analysis [RJE05] cannot cover all effects that modifications of one system property may have on system performance. One reason is that the variation of one property can also affect the values of other system properties requiring new approaches to keep track of simultaneous parameter changes.

Therefore, TU Braunschweig developed a heuristic and a stochastic approach for multi-dimensional sensitivity analysis [RHE06]. The heuristic approach is a divide-and-conquer like algorithm, which uses parameter specific heuristics to prune the search space. It is applicable to two dimensional search spaces. The stochastic approach is based on evolutionary search spaces and uses tabu search to bound the region containing the sought-after sensitivity front

separating working and non working system configurations. It is applicable to search spaces of arbitrary dimension.

MPARM interface with Lisatek (University of Bologna)

New processor models have been included. The most important extension in this area is the integration with the SystemC models generated by the Lisatek suite developed by AACHEN. Any processor modeled in LISA can now be integrated as add-on core in the MARM platform. A standardized transaction-level interface has been defined for core embedding.

MPARM memory models (University of Bologna)

Models for external memory controllers (DRAM-DDRAM). The main memory interface is often the true performance bottleneck for many MPSoC platforms. Therefore significant effort has been devoted to the development of an accurate DRAM controller module, capable of several advanced communication-optimizations. The model has been integrated within the MPARM platform. Associate partner STmicroelectronics has provided the functional specification for the controller.

Traffic generator model (University of Bologna, Technical Univeristy of Denmark)

Applications running on MPSoC architectures increasingly present non-trivial execution flows and synchronization patterns, especially in presence of underlying operating systems and when exploiting interrupt facilities. These properties make it very difficult to generate realistic test traffic. Technical University of Denmark and University of Bologna have jointly developed a reactive traffic generator device capable of correctly replicating complex software behaviours in the MPSoC design phase. The approach has been validated by showing cycle-accurate reproduction of a previously traced application flow. The traffic models have been integrated in both the ARTS environment from Technical University of Denmark and the MPARM environment from University of Bologna.

ARTS modelling framework (Technical University of Denmark)

ARTS is a SystemC-based abstract system-level modelling and simulation framework, which allows MPSoC designers to model and analyze the different *layers*, i.e., application software, middleware and platform architecture, and their interaction prior to implementation. In particular, ARTS provides a simulation engine that captures *cross-layer* properties, such as the impact of OS scheduling policies on memory and communication performance, or of communication topology and protocol on deadline misses. The ARTS framework was demonstrated at the University Booth at the DATE07 conference in Munich. As a result, ARTS has been made public available. The distribution consists of the framework and a tutorial. The results of this work was published at MASCOTS05 [MSM05] and an article has been submitted for the journal on Design Automation for Embedded Systems.

A web-link to the downloadable ARTS framework is http://www.imm.dtu.dk/arts .

Toolbox for Modular Performance Analysis method of ETHZ (ETH Zurich)

The analytic performance analysis model for distributed embedded systems and multiprocessing devices has been refined and discussed together with other partners. A major event has been the Distributed Embedded Systems workshop in Leiden and the Execution Platform Meeting in Bologna. As a result, we decided to implement the basic mathematical tools of Real-Time Calculus in form of a Matlab toolbox. The aim is to foster even more integration in the future as now other groups will be able to apply and incorporate analytic

methods easily. The first version of the toolbox is available, including documentation and a tutorial.

It will be used to integrate Symta/S and the modular performance analysis method in the next year of ARTIST2. A web-link to the toolbox is http://www.mpa.ethz.ch/Rtctoolbox/Overview .

Combining simulation and formal analysis for performance analysis (ETH Zurich)
Collaboration with Francesco Poletti and Luca Benini at University of Bologna

In this activity, we developed a new, compositional performance evaluation method for embedded systems. The new method combines existing approaches for system-level performance analysis, namely MPA a formal method and MPSim a simulation-based approach. To enable this combination, we defined the interfaces needed between the different performance evaluation methods. As a core of the approach, we propose a method to generate simulation stimuli from analytical models. In addition, we introduced a measure to assess the quality of a generated simulation trace with respect to its analytical description. In order to show the applicability of this new approach for performance evaluation, we implemented an example system for such a combined performance evaluation consisting of a multiprocessor system-on-a-chip. It is based on existing models for simulation and analytical models extended by the needed interfaces for the combination, including an implementation of the simulation trace generation algorithm. This combined model was then used for a case study of an application running on a multiprocessor system.

To achieve the results described above, several physical and phone meetings were held to coordinate the joint effort and to discuss future directions of this activity. The following two publications [KBPT06] and [KT06] describe the results of the joint activity.

Work in Year 3

NoC Emulation Framework (UNIBO, EPFL)

NoC-based MPSoCs involve new and critical design challenges, such as the design of network interfaces and protocols to provide reliable on-chip communication to transport the data of the cores. Also, the selection of suitable custom topologies of switches for the applications of the target MPSoC is critical to provide the needed low-latency at the physical interconnection layer to transport the data of the cores. All these challenges require a very time-consuming and error-prone design and tuning process of on-chip interconnects to design power-efficient and high-performance MPSoC.

UNIBO developed, in cooperation with EPFL a combined hardware-software NoC emulation framework, which shows how flexible NoC emulation can be used as a powerful design tool for tuning and functional validation of on-chip interconnections for MPSoCs. This emulation framework is implemented onto a *Field Programmable Gate Array* (FPGA) platform and has as one of its main novelties the utilization of the FPGA as an active element in the emulation control layer to speed up functional validation and to add flexibility to the NoC configuration exploration, instead of merely being the platform where the circuit is prototyped, as emulation is typically used.

The emulation framework is able to test actual physical realizations of NoCs on silicon up to four orders of magnitude faster than *Hardware Description Language* (HDL) simulators (see Figure 1), while preserving cycle accuracy. In addition, the flexibility of the emulation framework can be exploited to define a procedure to rapidly validate and tune NoC physical implementation characteristics (e.g., buffer size, topology of switches, size of inter-switches links, etc.) for real-life traffic patterns of software applications that can be executed in the target

MPSoCs or various software scenarios (e.g., bursts lengths, average on chip communication load, etc).

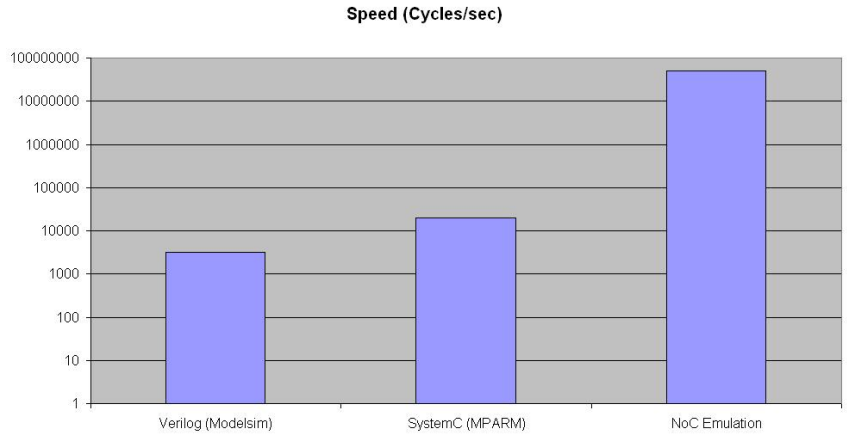

Figure 1: Speed comparison NoC-MPSoC simulation vs. emulation

Modeling and formal timing analysis of Multiprocessor Systems On Chip (TU Braunschweig)

TUBS have continued to investigate design paradigms of MPSoC architectures. As opposed to distributed systems, a common feature here is the use of a shared memory that is concurrently accessed from each processor, introducing conflicts on the memory and interconnects. System designers often implement latency-hiding techniques to reduce the effect of waiting for data, by allowing frequent context switches to tasks that are ready.

Building on previous work [SIE06] we have systematically investigated a realistic application with STMicroelectronics as an industrial partner.

The application was developed at the École Polytechnique de Montreal to run on the StepNP research platform. The involved round-robin scheduler could easily be integrated into our analysis engine. By conservatively considering the memory and bus congestion, this allowed to quickly model different architectural scenarios, and to predict corner case behavior which could not be identified in a simulation.

The work on the coupling a Synchronous Dataflow Graph based analysis with our compositional analysis approach has been presented at the DATE 2007 [SSE07]. It was received with great interest and has led to further cooperation with NXP Semiconductors, Eindhoven, NL.

Sensitivity Analysis and System Robustness Optimization for Complex Embedded Systems (TU Braunschweig)

TUBS have further extended their methods for sensitivity analysis and system robustness optimization.

As a result of HW/SW reuse, design data refinement or integration of components provided by different suppliers, the system designer must take into account that system properties, such as worst-case execution times, data rates, CPU clock rates, etc., are likely to be modified during the design process, or even later, during system life-cycle.

Our sensitivity analysis approaches can be used to compute, for the system properties subject to modification, the available slack with respect to an imposed set of constraints. Hence, any property modification carried out within the available slack interval guarantees that system

feasibility is preserved. In many cases, the modification of a system property implies also the variation of other properties in the system. For such cases, we developed a multi-dimensional sensitivity analysis [RHE06].

In order to efficiently control performance and to ensure predictability, sensitivity analysis must be systematically integrated into the design flow of embedded systems. We, therefore, proposed expressive robustness metrics for different assumptions and design scenarios, and showed how they can be efficiently considered throughout the whole design process. The proposed metrics are based on sensitivity analysis. At top level we distinguish robustness metrics for independent [HRE06] and dependent system properties [HRE07] w.r.t. system performance. For independent system properties the value of one system property does not have any influence on the admissible values for the other system properties. Contrarily, for dependent system properties the modification of one system property leads to more restrictions for the other system properties, i.e. their flexibility w.r.t. modifications decreases.

Performance characterization and system robustness become even more important if we assume that for complex application structures and dynamic scheduling policies, performance metrics, such as end-to-end latencies, response times, buffer sizes, etc., can easily exhibit unexpected non-monotonic behavior, a phenomenon known in literature as scheduling anomaly. In order to effectively cover such effects, we proposed a detailed scheduling anomaly analysis [RE06]. Our analysis can be used to find, on the one hand, system configurations with little design robustness, and on the other hand, to reveal additional performance reserves.

Fault-Tolerent Process Graph Model (DTU, LiU)

There is a lot of research in the area of system modeling and specification, and an impressive number of representations have been proposed. The system-level design tasks typically deal with sets of interacting processes. A process is a sequence of computations (corresponding to several building blocks in a programming language) which starts when all its inputs are available. When it finishes executing, the process produces its output values.

Researchers have used, for example, dataflow process networks (also called task graphs, or process graphs) to describe interacting processes, and have represented those using directed acyclic graphs, where a node is a process and the directed arcs are dependencies between processes. One drawback of dataflow process graphs is that they are not suitable to capture the different fault scenarios that can happen due to the occurrence of transient faults in a fault-tolerant application. For example, it can happen that the execution of some processes fails due to faults. By explicitly capturing such a failure in the process graph model, a more fine-tuned modeling and a tighter (less pessimistic) assignment of execution times to processes is possible, compared to traditional data-flow based approaches.

Together with Linkoeping University (LiU) DTU have proposed an extension to the process graph model, namely a “fault-tolerant process graph” model (FT-PG). In an FT-PG the fault occurrence information is represented as conditional edges, and thus the FT-PG captures all the fault scenarios that can happen during the execution of application [TVLSI]. We have shown how design transformations that introduce redundancy, such as re-execution and replication, can be applied on this model.

MOVES, Modeling and Verification of Embedded Systems (DTU, AAU)

One of the major challenges in designing an embedded system is to find a mapping of the application onto the execution platform, which effectively fulfills the non-functional requirements of the embedded system such as timing, memory usage, energy consumption, and other cost. A particular challenge is to model and analyse cross-layer dependencies, where the change of a property in one part of the system, e.g. scheduling policy, may impact the performance of another part of the system, e.g. deadline miss on another processor, and hence, the overall

system performance. The ARTS simulation model developed by DTU during the first two years of ARTIST2, has been modeled using the semantics of timed automata and implemented in UPPAAL from AAU.

In order to make the formal model available for easy adaptation of embedded systems designers, the UPPAAL based model has been embedded in a tool called MOVES. MOVES supports formal analysis of non-functional properties of an embedded system, covering the system layers of an application mapped on an execution platform, consisting of a heterogeneous multiprocessor architecture where each processor may run a real-time operating system, and where all processors are connected through a network. It supports the designer by allowing him/her to describe the application, the execution platfrom and the mapping in a straight forward manner. MOVES then translates the system into a UPPAAL model which is then used to model check the system against given properties. If the model checking fails, the given counterexample produced by UPPAAL, is translated by MOVES into a schedule indicating where the properties were violated. The designer can then use this information to understand why the system failed and to suggest improvements.

Modeling and Verification of Hardware Components (DTU)

As the complexity of chips grows, the methodology to build chips has to evolve. Today, chips are largely synthesized from high-level architectural descriptions that hide low-level details.
The majority of hardware designs are done using the most common hardware description languages, suvh as VHDL or Verilog. Both languages support high-level architectural descriptions, but allow hardware designers to incorporate low-level details in order to optimize for a particular hardware technology and directly synthesize using a restricted subset of the languages. However, chips may also be synthesized from software based models in much the same way as compilers produce executable code. Examples of such languages are Esterel, Lustre and Signal.
DTU have developed a language for hardware models based on the Gezel hardware description language developed and maintained by Virginia Tech, USA. The language depends on reasonably few, simple and clean concepts, and it strikes a balance between software and hardware concerns that suits the needs for a modern top-down approach to hardware design.

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

Simulation Platform for Dynamical Reconfigurable Systems (DTU)

One of the biggest challenges in reconfigurable system design is to improve the rate of reconfiguration at run-time by reducing the reconfiguration overhead. Such overhead comes from multiple sources, and without proper management, the flexibility of the reconfiguration can not justify the overhead cost. DTU have developed a flexible framework, called COSMOS, to model and simulate coprocessor-coupled reconfigurable systems. The framework is an extension to the ARTS framework developed by DTU during the first two years of ARTIST2. DTU have develpoed a novel real-time task model that captures the characteristics of dynamically reconfigurable systems' task in terms of initialization, reconfiguration and reallocation. DTU also propose a general model of coprocessor-coupled reconfigurable systems. The task and architecture models have been extended to facilitate the study of run-time resource management strategies. Based on this model, DTU have demonstrated how a simple "worst case" run-time system can be modelled in the COSMOS framework as a firmware to manage the application execution.

The COSMOS framework have been used to experiment with various combinations between the application and the architecture to gain a better understanding of the critical issues in reconfigurable architecture design. A set of experiments based on a MP3 task graph have been conducted.

MT-ADRES: Multithreading on Coarse-Grained Reconfigurable Architecture (DTU, IMEC)

To investigate the performance bottleneck and the scalability of the state-of-the-art datapath-coupled reconfigurable architectures, DTU and IMEC have studied the coarse-grained reconfigurable architecture ADRES (Architecture for Dynamically Reconfigurable Embedded Systems) developed by IMEC, Belgium. In order to improve task-level parallelism, they have proposed a method for multithreading on the ADRES architecture.

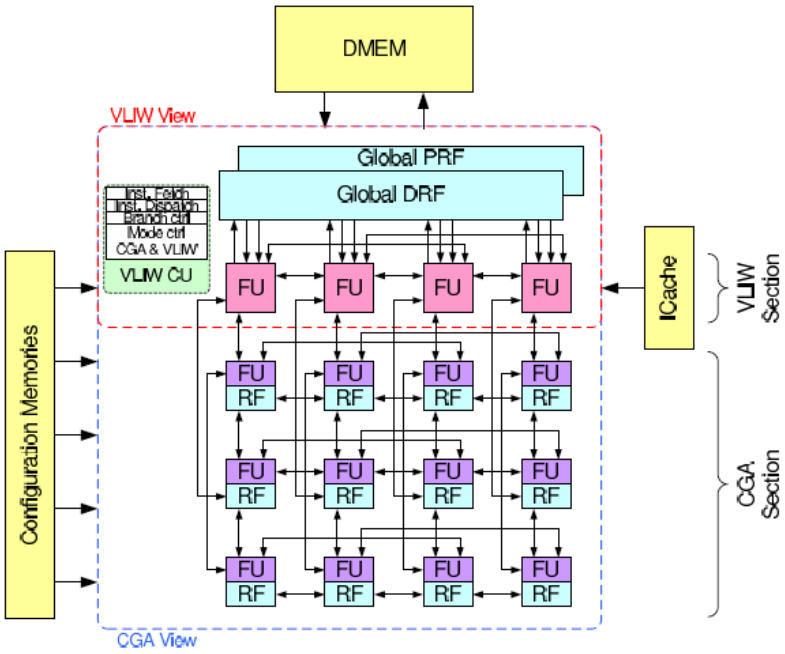

DTU and IMEC have proposed how the ADRES architecture can be extended to support multi-threading, and how the ADRES compilation tool flow needed to be extended to cope with multithreading. They have made an experiment running a dual-threaded MPEG2 decoder on a customized ADRES architecture to demonstrate that multi-threading is feasible for ADRES. Through the MPEG2 experiment they have discovered some design pitfalls that hinder the performance of the threaded ADRES, and discussed what technologies can further improve the performance of the multi-threaded ADRES.

MPA Model Integration (ETHZ, TUBS)

ETHZ has been combining its tool set (MPA – Modular Performance Analysis) with the system modelling infrastructure of other partners. Especially, there has been a deep integration between Symta/S and MPA. This not only entails converters between the two modeling formalisms but also investigations, when to use which formalism. In particular, it turns out that there are components of a design that can be much more accurately modeled by one or the other model. Besides the tool integration, there has been a paper published at CODES/ISSS that describes the obtained results.

In addition, there was an integration of the PISA multi-objective optimization framework into Symta/S as well as in tool to determine the robustness of a design at TU Braunschweig. Therefore, the exchange of tools and the integration between different modeling formalisms

and too domains has been succesfully demonstrated. These goals heve been achieved by means of mutual visits, e.g. Simon Kuenzli (ETHZ) spend time at TU Braunschweig in October 2006.

Simulation platform for distributed embedded systems (LiU)

LiU have finalized their simulation platform for distributed embedded systems. Once the platform was available the efforts were concentrated into the following two directions:

1. Elaboration of a simulation methodology which allows to efficiently estimate the worst case response time of distributed real-time applications. In order to achieve an efficient simulation, two problems had to be solved:
   a) how to reduce the space of execution times to be explored;
   b) how to generate the next exploration point at a given moment of the simulation process? In other words, what exploration strategy to use.

   Simulations, if well conducted, can lead to tight lower bounds on worst-case response times, which can be an essential input at design time. Moreover, such a simulation methodology is very important in situations when the running application or the underlying platform is such that no formal timing analysis is available.
2. LiU have used the elaborated simulation platform two validate formal analysis approaches, by estimating their degree of pessimism. They have performed such an estimation of pessimism on two response-time analysis approaches for distributed embedded systems based on two of the most important automotive communication protocols: CAN and FlexRay.

Modeling and analysis for NoC communication (LiU)

The Linköping group has continued its work on the modeling and analysis of NoC platforms. In particular fault tolerance in the context of NoCs and transient faults has been addressed. Based on the elaborated modeling and analysis approach a system optimization methodology has been developed.

Final Results

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

The focus for year 4 was mostly to continue and extend the formal-based models, although some work on simulation-based models would be continued. For the formal-based models, the focus was to continue and extend the semantic model of SymTA/S to efficiently cover MPSoC architectures and to further investigate challenging timing issues in multiprocessor systems. The work on models for the analysis and optimization of fault-tolerant embedded systems should be continued and further extended to capture fine-grained combinations of several fault-tolerance techniques. Two major efforts were planed for the models based on timied automata; 1) to continue the work on formalizing the ARTS model using timed automata based on UPPAAL, and in particular, to refine the formal model to address modeling and verification issues closer to the hardware layer of the execution platform. 2) to combine Modular Performance Analysis with timed automata based evaluation methods.
For the simulation-based models, the focus was to continue the work on extending the simulation-based model towards handling dynamically reconfigurable architectures, and iIn

particular, to study different run-time resource management strategies. And for the simulation environment for distributed embedded systems, to do experimental evaluations to address the level of pessimism obtained through the formal methods.

The objectives have all been achieved. In the following, details are given for each sub-activity, listing a title and the partners involved in the sub-activity.

Modelling and Formal Performance Analysis of Multiprocessor Systems On Chip (TU Braunschweig)

TU Braunschweig has continued to investigate design paradigms of MPSoC architectures. As opposed to distributed systems, a common feature here is the use of a shared memory that is concurrently accessed from each processor, introducing conflicts on the memory and interconnects.

We have additionally investigated the options of applying the optimization techniques known from distributed multiprocessor systems with the analysis methodologies for multiprocessor systems on chip. The approaches turned out to be flexible enough to be easily adapted to exploring parameters of memory controllers and traffic shapers. With this, an optimal configuration for a realistic media processing application could be found, in which a careful balance between worst case latency and guaranteed throughput is required from the memory subsystem. The results of this experiment have been presented in [SHR+08].

With the investigation of a realistic application by the École Polytechnique de Montréal and STMicroelectronics as the supplier of the hardware models, we could also benchmark our analysis approach. Our conservative models have turned out to overestimate the worst simulated values by no more than 25% in the given system, which is a fair value given the increase in design confidence. The results will be published in [SNN+08].

Simulation-based and analytical methods for validation of distributed real-time systems (Linkoping)

We have finalized and improved our simulation-based and analytical platforms for validation of distributed embedded systems. One of the main directions was the evaluation of the simulation-based platform with regard to its capacity to be used for validation of time critical systems. Based on experimental results with the simulation-based platforms we were also able to draw interesting conclusions regarding the potential pessimism of formal approaches to validation of distributed real-time systems. The experiments were performed considering FlexRay and CAN-based distributed systems, which are of great interest for the automotive industry. Results were published at DATE 2008.

Integration of thermal models in a framework for energy efficient design of time constrained embedded systems (Linkoping)

This work is built on top of our research regarding energy efficient design of real-time systems. We have earlier developed several approaches to dynamic voltage selection (DVS), a technique which exploits the available slack times by reducing the voltage and frequency at which the processors operate and, thus, achieves energy efficiency.

It is known that high power densities achieved in current SoCs do not only result in huge energy consumption but also lead to increased chip temperatures. Growing temperature leads, among others, to an increase in leakage power and, consequently, energy, which, again, produces higher temperatures. Nevertheless, the temperature issue has been completely ignored in the DVS techniques for real-time embedded systems proposed in literature.

We have used state of the art temperature models for multicore chips and integrated them into a temperature aware dynamic voltage selection approach. We have shown that important energy savings can be achieved. Results have been published at DATE 2008.

Modeling and analysis of fault tolerant distributed embedded systems (LiU, DTU)

LiU and DTU have continued their cooperation in the area of fault-tolerant real-time systems. The main focus was extending the system model in order to capture not only hard real-time applications but also soft real-time systems and, more generally, systems consisting of both hard and soft real-time tasks. The goal is to guarantee deadlines for the hard processes even in the case of faults, while maximizing the overall utility. We use time/utility functions to capture the utility of soft processes. Results have been published at DATE 2008.

DTU and LiU have extended their fault-tolerant process graph model for embedded systems to consider a combination of hardware and software fault-tolerant techniques. The model can capture re-execution of processes and hardened, i.e., more reliable, hardware processing elements. A System Reliability Analysis (SRA) technique has been proposed that calculates the reliability of the system considering: (a) the different hardening levels in hardware, (b) the re-execution levels in software (how many redundancies are there for a process), (c) the mapping, which decides what hardening-level will a process start from and (d) the scheduling, which decides how are the recovery slacks shared. Proposing such a SRA has been challenging because, due to the slack sharing, there are many complex combinations in which an application can fail.

MOVES, Modeling and Verification of Embedded Systems (DTU, AAU)

DTU has continued their work on formalizing the ARTS simulation model and to make it usable for designers early in the design process. In order to support designers of industrial applications, the timed-automata model is hidden for the user, allowing the designer to work directly with the abstract system-level model of embedded systems. The designer provides an application consisting of a set of task graphs, an execution platform consisting of processing elements interconnected by a network and a mapping of tasks to processing elements. The system model is then translated into a timed-automata model which enables schedulability analysis as well as being able to verify that memory usage and power consumption are within certain limits. In the case where a system is not schedulable, the tool provides useful information about what caused the missed deadline. DTU does not propose any particular methodology for design space exploration, but provide an analysis framework, MoVES, where embedded systems can be modelled and verified in the early stages in the design process. Thus, the MoVES Analysis Framework provides tool support for system designers to explore alternatives in an easy and efficient manner.
An important aspect in the design of MoVES is to provide an experimental framework, supporting easy adaptability of the ”core-model” to capture energy and memory considerations, for example, or to experiment with, say new principles for task scheduling and allocation. Furthermore, the MoVES Analysis Framework is equipped with different underlying UPPAAL models (some of which have been developed together with AAU), aiming at efficient verification in various situations. For the moment DTU is operating with the following underlying models for

- schedulablity analysis in connection with worst-cases execution times only,
- schedulablity analysis for the full core model (including best- and worst-case executing times),
- schedulability analysis addressing memory and energy issues as well, and
- schedulability analysis for the full core model on the basis of stop-watch automata. This

analysis approach is based on over approximations, but it has provided exact results in the experiments carried out so far and it appears to be the most efficient Uppaal implementation.

Formal verification of design properties of hardware architectures (DTU)

DTU has continued its work on a formal language for hardware models based on the Gezel hardware description language developed and maintained by Virginia Tech, USA. The language depends on reasonably few, simple and clean concepts, and it strikes a balance between software and hardware concerns that suits the needs for a modern top-down approach to hardware design. The semantical domain of the language has been related to timed-automata using the UPPAAL system. They have demonstrated a formal verification of design properties of a few simple example circuits including the Simplified Data Encryption Standard (SDES) Algorithm and different algorithmic implementations of the Greatest Common Divisor. Verification guarantees properties of the underlying algorithm, e.g. correct output for any given input, as well as other properties such as upper limits on the number of clock cycles for the algorithm to stabilize with a given input and upper limits on the number of register updates, to serve as an indicator of energy consumption.

Simulation Platform for Dynamical Reconfigurable Systems (DTU)

Understanding the dynamic behavior of run-time reconfigurable systems is a very complicated task, due to the often very complicated interplay between the application, the application mapping, and the underlying hardware architecture. However, it is a key issue to determine the right reconfigurable architecture and a matching optimal on-line resource management policy. Although architecture selection, application mapping and run-time system have been studied intensively in the past, they have not been thoroughly studied and modelled in the context of run-time reconfigurable system. DTU has extended its simulation framework COSMOS to study the dynamic behavior of run-time reconfigurable systems. Through a number of design space exploration experiments, they have pinpointed the critical design issues in the reconfigurable architecture study and analyze their impact on the architecture performance.

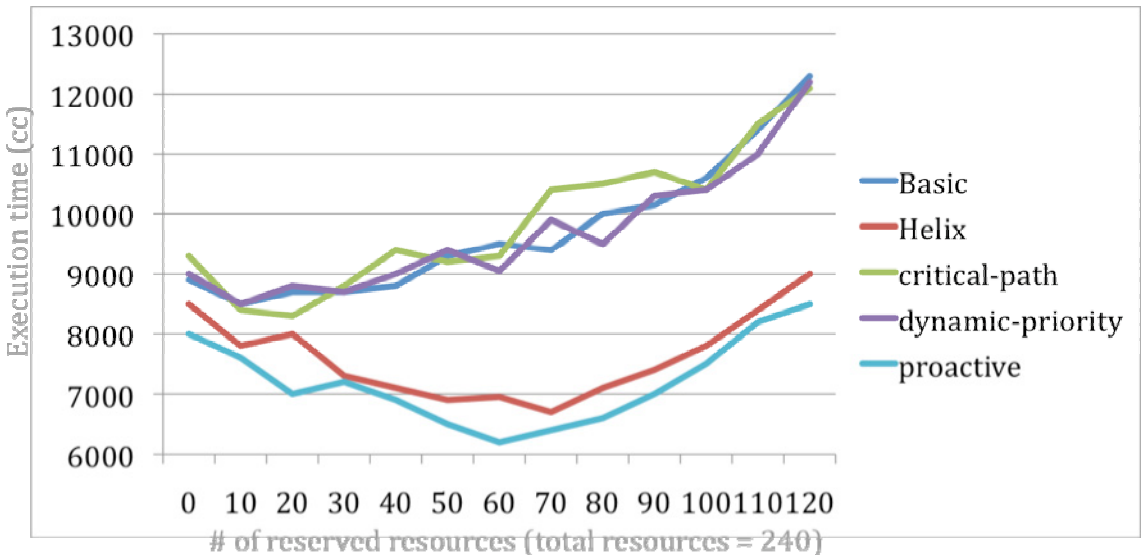

Analysis tools for embedded systems (DTU, Oldenburg)

This activity addresses the problem of establishing decidability and model-checking results for fragments of interval logics. The aim is, in particular, to establish efficient tools for the analysis

of real-time properties of embedded systems, and, in general, tool support for the analysis of more general resource constraints of embedded systems. The activity is a co-operation with Prof. Martin Fränzle, Oldenburg University.

Composing analysis frameworks

As proposed, ETH Zurich coupled two different analysis frameworks, i.e. MPA (Modular Performance Analysis) and Timed Automata. This work has been done with the affiliated partner NUS (P.S. Thiagarajan). To this end, one of his PhD students visited ETH for a duration of 6 months. The results have been published and are now the basis for further work. In addition, ETH Zurich has been leading an activity where several partners from ARTIST2 have been involved which compared various abstraction mechanisms used in the analysis of distributed embedded systems. This work resulted in a joined conference and in a joined journal publication.

Modeling patterns for performance analysis (TU/e, ETHZ)

Modeling patterns for specification of applications and platforms and the automatic transformation to executable models for performance analysis. The key idea is to decouple the application and platform specification from the specific formal modeling and analysis tools. (TU/e work)

Comparison of different Y-chart based approaches (Metropolis, MPA and SHE/POOSL) for design-space exploration (combined work of University of Montral (Canada), TU/e and ETH Zurich)

Resource-aware Design (NoE Integration)
Led by Luca Benini (University of Bologna)

Partner teams (leaders): Luca Benini - University of Bologna (Italy), Petru Eles - University Linköping (Sweden), Rainer Leupers - RWTH Aachen (Germany), Jan Madsen, Technical University of Denmark (TUD), Peter Marwedel - University of Dortmund (Germany), Lothar Thiele – ETH Zürich (Switzerland), Reinhard Wilhelm – Saarland University (Germany)

Affiliated teams (leaders): Roberto Zafalon – STM (Italy)

Overview: The importance of resource awareness in embedded systems growing. Resources include energy, computational power and hardware components. This activity is concerned with optimization of resource usage, while at the same time meeting application requirements. Energy efficiency, Reliability are typical requirement of embedded applications. Predictability is also an extremely important one.

With the growing software content in embedded systems and the diffusion of highly programmable and re-configurable platforms, software is given an unprecedented degree of control on resource utilization. This relation between hardware and software layers can be used to perform aggressive optimizations that can be achieved only by a synergistic approach that combines the advantages of static and dynamic techniques.

Work in Year 1

-- Work achieved in the first 6 months: --

Cooperation was established between the Universities of Bologna and Dortmund. The objective was to integrate the memory-aware compiler developed in Dortmund with the multi-processor platform simulator developed in Bologna. The first results were the definition of a

standard format for the executable output of the compiler, as well as for the memory allocation information. This output format is supported by the platform simulator.

Cooperation was furthermore established between the Universities of Bologna and Aachen. The objective is to extend the modelling capabilities of the platform simulator developed in Bologna toward heterogeneous multi-core architectures, exploiting the Application-specific Processor development framework based on the LISA architecture description language developed in Aachen. The first result of this work was the definition of a standardized wrapping protocol which allows any SystemC core description generated by Aachen tools to be instantiated (multiple times) as a core in the platform simulator by Bologna.

-- Work achieved in months 6-12 --

The cooperations established in the first six months were continued and were significantly strengthened, as a significant amount of technical work was performed to sustain them. More specifically:

- The cooperation between Bologna and Dortmund required the development of a new souce-level transformation tool for performing memory optimizations by Dortmund, and the development of compatible memory organization models by Bologna (including I and D caches as well as scratchpad memories).
- The cooperation between Aachen and Bologna required extensive re-design of Bologna's core interfacing protocol within the platform simulator. On the other hand, Aachen has provided extensive technical support on Lisatek core wrapping architectures and toolsets.

An additional cooperation between Bologna and Saarland University was established. The objective of this cooperation is the exploitation of the platform simulator developed at Bologna, more specifically of the timing accurate core models incorporated in the simulator, as targets for the worst case execution analysis framework developed in Saarland University.

Work in Year 2

Sorting of the following contributions is by the name of the city of the partner. Partners are mentioned in alphabetical order of the cities. The sequence of partner institutions has no significance with respect to the share of the partners in the workload.

Integration of LISATek ISS models in SystemC and the MPARM virtual platform (Aachen, Bologna)

The issues arising from the integration of LISATek ISS models in SystemC and the MPARM virtual platform have been investigated in more detail [Ang05], especially concerning the interaction with level one (L1) memories. A new MPARM functional model was developed to handle the L1 memory. It was also useful to cluster other functionality within the same block. The end result is called a "processor tile", comprising LISATek-generated SystemC model of the processor and the most tightly coupled components (see fig. 1).

The following component models were developed:

- a timer device,
- an emulated serial port,
- a simple interrupt controller.

The first component is vital if attempting to port an operating system. The second is very useful for debugging purposes; placing it next to IP cores, instead of in a shared location accessible

to all system processors, has the advantage of allowing for independent input/output, and prevents debug traffic from spilling onto the system interconnect where it could pollute performance statistics. Finally, the interrupt controller is both a requirement of the other two devices and a crucial component to develop efficient synchronization mechanisms in multiprocessor systems. The controller is externally attached to a set of system-level wires which convey inter-core interrupts. On the IP core side, a simple interrupt handshaking protocol was implemented at Bologna. In this protocol, the value of interrupt registers is copied on some LISATek core pins which are polled every cycle by the core to take proper action. The interrupt controller is memory mapped to let the core reset the pending interrupt flags and configure the masking status.

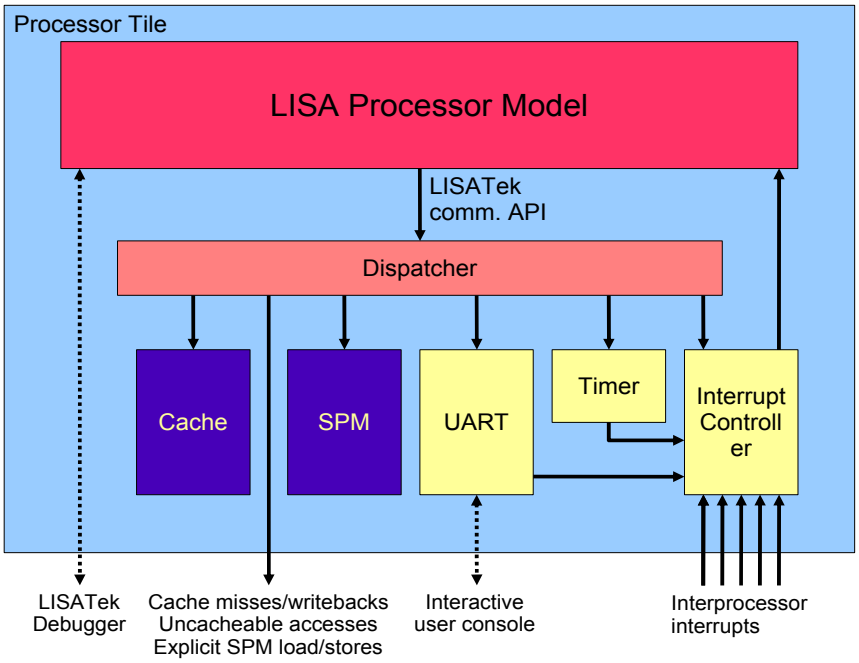

Figure 1 Processor Tile

Energy efficient time constrained systems (Bologna, Linköping)

Power models as well as a simulation environment for validation have resulted from cooperation of the University of Linköping with the Bologna group. As the first step, an approach for mono-processor systems has been elaborated, implemented and published [And05].

During the last six months of year 2, the efforts concentrated on an extension of this approach to multiprocessor systems. During the summer 2006, this work was performed as part of the ARTIST mobility action in cooperation with the Dortmund group and extended into the reporting period of year 3.

Predictability in Multiprocessor System on a Chip (MPSoC) architectures (Bologna, Braunschweig, Linköping)

Besides being energy efficient and having a high performance, for many applications it is required that multiprocessor SoC implementations are highly predictable with respect to their timing behaviour. This problem has been addressed by the Linköping group during this period. While this issue has been previously investigated in the context of mono-processor systems, available results are inapplicable to modern multiprocessor architectures in which, for example,

due to the shared memory access and shared buses, the individual WCETs of tasks depend on the global system schedule. Providing WCET guarantees and reliable schedules in this context is extremely challenging. It involves issues related to bus protocols and control, WCET analysis, system level scheduling and optimizations. With regard to the "classical" aspect of WCET analysis the group is building on the Symta/P tool from the Braunschweig group (Ernst et al., a member of the (“execution platform” cluster). The Linköping group is also interacting with the Bologna group with regard to the issues of bus control.

This work is an effort started at the beginning of 2006. The overall concept has been elaborated, solutions have been developed and tools are under implementation. Publications and further results are expected in the following period.

Web site: http://www.ida.liu.se/~eslab/real-time.html

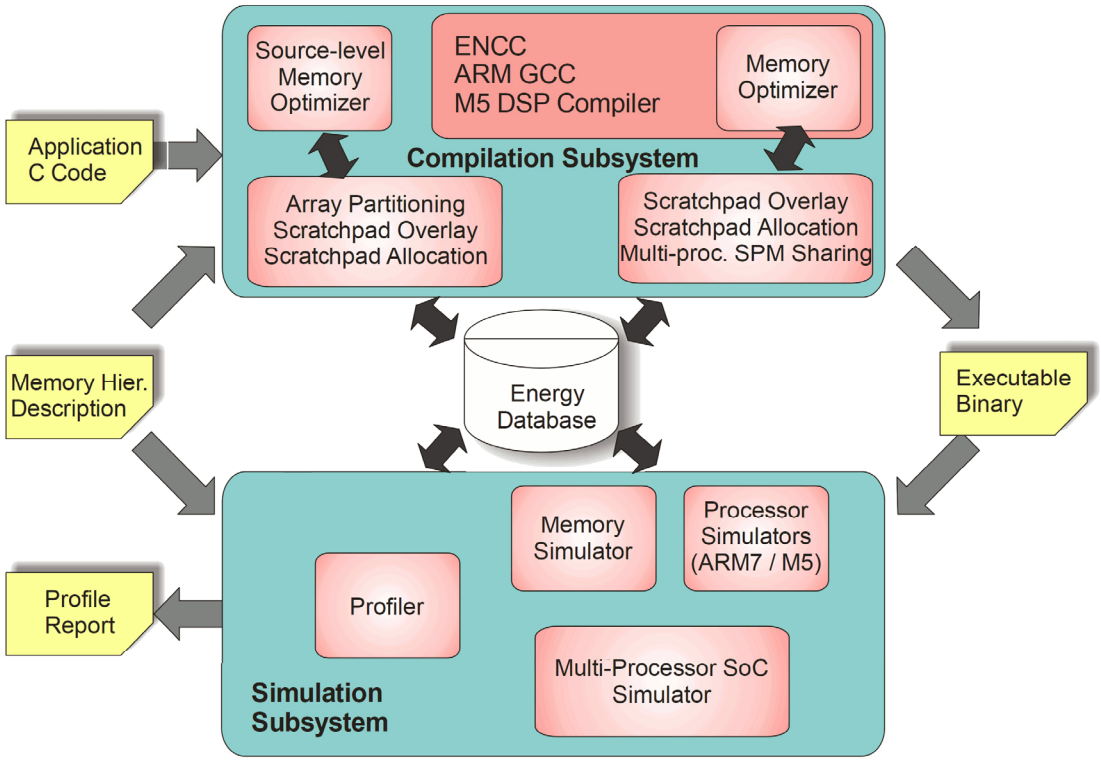

Figure 2: Memory Aware Compilation and Simulation Tool-Chain

Memory Aware Compilation and Simulation Tool-Chain for Energy Optimizations (Bologna, Dortmund)

During the last reporting period, the need for a coherent tool chain for energy optimizations and for exploration of memory hierarchies across different system architectures was recognized. Therefore, a memory aware tool-chain supporting uni-processor ARM, multiprocessor ARM and M5 DSP based systems was developed (see fig. 2) at Dortmund. Both the simulation and compilation subsystems are configured from a single memory hierarchy description. In addition, a common energy database is used by the memory optimizers in the compilation subsystem as well as by the memory and multi-processor SoC simulators in the simulation subsystem. The developed tool-chain optimizes input application code for a given memory hierarchy [Ver06d, Weh06] and also evaluates the optimization by simulating the optimized executable on the same memory hierarchy. The tool-chain is developed due to the cooperation

between University of Dortmund and University of Bologna, as the simulation subsystem includes the multi-processor SoC simulation from Bologna while the compilation subsystem is developed at Dortmund. Moreover, both partners have agreed on a common memory hierarchy description format, which will be used for developing future optimizations [Ver06c].

Web site: http://ls12-www.cs.uni-dortmund.de/research/macc

Design-Time Memory Allocation Techniques for Multi-Process Applications with Aperiodic Processes (Bologna, Dortmund)

Previous work at Dortmund proposed compile-time or design-time memory allocation approaches to share the scratchpad memory among the periodic processes of a multi-process application. The current work extends the previous work and proposes memory allocation approaches for applications consisting of aperiodic tasks. This significantly increases the complexity of the memory allocator as the arrival times of the processes are completely unknown at design time. Therefore, the memory allocator is divided into an intelligent design-time component and a simple run-time component.

The design-time component of the memory allocator works in the following stepwise manner. First, it identifies memory objects, *i.e.* code segments and data variables, which on scratchpad allocation lead to reduction in the energy consumption of the system. Second, it processes the application code to enable the movement of memory objects at runtime. Finally, it inserts blocking statements in the application code to prevent unsafe movement of memory objects. The runtime component, depending upon the current set of active processes and the current state of the scheduled process, allocates (de-allocates) memory objects to (from) the scratchpad memory. Experiments report that a two-phased memory allocator minimizes the energy consumption due to applications with aperiodic tasks [Ver06a, Ver06b].

Web site: http://ls12-www.cs.uni-dortmund.de/research/macc

Operating System Support for Online Allocation of Scratchpad Memories (Bologna, Dortmund)

The goal of this work at Dortmund is to develop a runtime memory allocator which keeps track of the execution behaviour of the application and allocates scratchpad memory with memory objects (code segments and data variables) at runtime. The runtime allocator of this approach is more complex than the design-time memory allocator described above. At compile time, attributes such as access counts and the size are computed for each memory object. These attributes are then supplied as input to the memory allocator. The allocator based upon the input attributes, the scratchpad memory utilization and the current execution pattern swaps memory objects in and out of the scratchpad memory. Several heuristics as well as analytical approaches have been proposed for the online allocation of the scratchpad memory. The proposed approaches have been integrated into the RTEMS operating system. Experiments demonstrate that for highly dynamic applications, significant energy savings can be achieved.

Web site: http://ls12-www.cs.uni-dortmund.de/research/macc

Resource awareness in sensor networks (Bologna, ETH Zürich)

The University of Bologna cooperated with ETH Zürich on resource awareness in sensor networks. For a full description please refer to the report on progress within the execution platform cluster for year 2.

Resource aware design space exploration (DTU, ETH Zürich)

The Technical University of Denmark (DTU) has developed a multi-objective design space exploration environment based on the PISA environment for multi-objective optimization from

the group of Lothar Thiele, ETH Zurich. The exploration is based on a genetic algorithm to solve the problem of mapping a set of task graphs onto a heterogeneous multiprocessor platform. The objective is to meet all real-time deadlines subject to minimizing system cost and power consumption, while staying within bounds on local memory sizes and interface buffer sizes. The approach allows for mapping onto a fixed platform or onto a flexible platform where architectural changes are explored during the mapping. This work will be continued. A paper was published at DIPES 2006 [Mad06].

FET Open Call project proposal (Dortmund, ETH Zürich, Saarbrücken)

A consortium from within ARTIST2 consisting of the Universities of Saarbrücken, Zürich, Bologna, Pisa and Dortmund as well as AbsInt has applied for a project on “Reconciling Performance with Predictability” in the FET Open Call. Both short and long proposals have passed all thresholds.

Analysis of cache predictability (Saarbrücken, AbsInt)

First quantitative results have been obtained on the predictability of different cache architectures. A paper is in preparation.

Improvement of timing analysis by integration with code synthesis (Saarbrücken, AbsInt)

The University of Saarbrücken and AbsInt (an industrial member of the compiler cluster) have cooperated with ETAS (an external company located at Stuttgart, see http://www.etas.com) on the integration of the ASCET-SD model-based design tool with the AbsInt timing analyzer aiT. This work is continuing. A paper was published by Ferdinand et al. [Fer06].
http://en.etasgroup.com/about/tradeshows/documents/2006-03-15_AutomotiveSoftwareWorkshop_ASCET_Paper_Renz.pdf

Interfaces for real-time components (ETH Zürich, EPFL)

Between members of the group of Tom Henzinger (EPFL) and Lothar Thiele (ETHZ) there have been intensive discussions on interface based design of embedded systems. There were common meetings and presentations. The main concept is to extend the common idea of static types towards resource types that talk about the use of various resources by a component, e.g. power, time, computing resources. As a result, the concept of interface-based design (by Tom Henzinger) has been successfully applied to real-time systems and associated publications have been written [Hen06, Thi06, Cha06]. http://chess.eecs.berkeley.edu/pubs/92.html

Work in Year 3

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

The very first steps for this work have been done during the first and second reporting period. The work has continued during the third year and first results are available and have been published.

In multiprocessor systems, the traffic on the bus does not solely originate from data transfers due to data dependencies between tasks, but is also affected by memory transfers as result of cache misses. This has a huge impact on worst-case execution time (WCET) analysis and, in general, on the predictability of real-time applications implemented on such systems. As opposed to the WCET analysis performed for a single processor system, where the cache miss penalty is considered constant, in a multiprocessor system each cache miss has a variable penalty, depending on the bus contention. This affects the tasks' WCET which, however, is

needed in order to perform system scheduling. At the same time, the WCET depends on the system schedule due to the bus interference. We have developed an approach to worst-case execution time analysis and system scheduling for real-time applications implemented on multiprocessor SoC architectures. An important aspect of the problem is the bus scheduling policy and its optimization, which is of huge importance for the performance of such a predictable multiprocessor application. What concerns the "classical" aspect of WCET analysis we are building on the Symta/P tool from the Braunschweig group. The design of appropriate bus controllers to support the proposed approach is done in cooperation with the group in Bologna. A master student from Bologna is visiting Linköping for a period of seven months starting with June 2007.

Energy efficient time constrained systems (Bologna, Dortmund, Linköping)

This work has been started in the previous reporting period and has been performed at Linköping in cooperation with the groups at Dortmund and Bologna.

Olivera Jovanovic, a master student from Dortmund visited Linkoping for 7 months. The work has aimed at extending a dynamic, on-line, voltage scaling approach so that it can be applied to multiprocessor systems. Another extension concerns taking into consideration the voltage/frequency switching overheads at energy optimization. An approach for dynamic and leakage energy reduction via combined supply voltage scaling and body biasing in real-time multiprocessor systems has been developed. Discrete voltage modes and intra-task scaling have also been considered. The mapping and scheduling of the task sets were assumed to be already given. The main optimization target for this approach is to achieve energy efficiency by exploiting dynamic slack, which results at runtime, for example when the tasks do not execute their worst case number of clock cycles. Dynamic slack is exploited by using online voltage reduction techniques. Since these algorithms are executed online, after each of the tasks finishes, they must have a low complexity. Moreover the energy and time overhead for changing the supply and body bias voltage is also considered. Results for a journal submission have been generated.

For the experimental validation of the approach the MPARM simulation platform from Bologna has been used. A publication reporting the research results is in the final refinement steps.

Fault-tolerant embedded systems design (DTU, Linköping)

The Technical University of Denmark (DTU) and Linköping University started a collaboration on safety-critical embedded systems. Safety-critical applications have to function correctly and meet their timing constraints even in the presence of faults. Such faults can be permanent (i.e., damaged microcontrollers or communication links), transient (e.g. caused by electromagnetic interference), or intermittent (appear and disappear repeatedly). The transient faults are the most common, and their number is increasing due to the increasing level of integration in semiconductors.

Linköping has proposed a list scheduling-based heuristic for the generation of fault-tolerant schedules, and have used a tabu-search meta-heuristic on top of list scheduling to optimize the assignment of fault-tolerance policies (i.e., re-execution vs. active replication) in order to reduce the fault-tolerance overheads. Such heuristics are able to produce good quality solutions in a reasonable time.

Researchers have used constraint logic programming (CLP) in the context of system-level design. The advantages of a CLP approach are: it produces optimal solutions, can capture complex design constraints and trade-offs, it is flexible, more general and easy to extend. However, none of the proposed CLP approaches take into account fault-tolerance aspects.

Hence, DTU has proposed a CLP framework that produces the fault-tolerant schedules such that the application is schedulable in the presence of transient faults, and the constraints and tradeoffs imposed by the designer are satisfied.

DTU have modelled the application as a fault-tolerant process graph, where the fault occurrence information is represented as conditional edges, and they have proposed an algorithm for the derivation of such graphs. The proposed CLP framework can be used to easily capture design optimization problems such as mapping and fault-tolerance policy assignment. In addition, the CLP framework can be used to reason about the effects of voltage scaling on reliability. Then, the system can be optimized for energy minimization under limited resources and strict timing and reliability constraints.

DTU have compared their CLP scheduling approach with the list-scheduling proposed by Linköping, and the CLP performs 25% better on average. By carefully optimizing the system implementation they are able to provide fault-tolerance under limited resources. The cooperation with Linköping has been in terms of reciprocal visits (Paul Pop, DTU has visited Linköping several times in 2007 and Viacheslav Izosimov, Linköping has visited DTU during 2006), exchange of tools and case studies.

Integration of LISATek ISS models in SystemC and the MPARM virtual platform (Aachen, Bologna)

The integration between the LISATek flow and the MPARM virtual platform has completed in year 3. As many embedded systems today deploy multiple instantiations of very-long-instruction-word (VLIW) processors for dataprocessing, integration efforts have aimed at developing a VLIW core suitable for multi-instantiation in the MPARM virtual platform. A VLIW architecuture compatible with the VEX instruction set (a simplified version of the STMicroelectronics' ST230 ISA) has been developed in LISA.

Significant effort has been devoted to exploiting the capabilitiy of the LISATek tools to generate synthesizable VHDL, if a suitable sub-set of the LISA syntax is utilized for the processor description. A 4-stage pipelined architecture has been described.

A SystemC model for virtual platform integration and VHDL model for synthesis have been automatically generated using LISATek tools. The VHDL version was synthesized from standard cells using Synopsys Design Compiler. In a first phase, UMC 0.13 µm technology with Design Compiler version 2004.12-SP2 was used and in a second phase there was a migration to a newer TSMC 90nm technology with the newer Synopsys Physical Compiler Y-2006.06.

Memory Aware Compilation and Simulation Tool-Chain for Energy Optimizations (Bologna, Dortmund)

In this activity Bologna has focused on developing a software infrastructure for compiler-based parallelization for MPSoC platforms. One of the key components of any compiler-parallelized code is barrier instructions which are used to perform global synchronization across parallel processors. As compared to programmer-parallelized codes, compiler-parallelized codes can contain larger number of barriers, mainly because a compiler has to be conservative in parallelizing an application (to preserve the original sequential semantics of the program), and this means, in most cases, inserting extra barrier instructions in the code.

Bologna has worked towards the implementation of MPSoC-suitable lightweight runtime synchronization facilities used by a parallelizing compiler frontend, with particular emphasis on barrier implementation. In order to avoid overheads due to multiple software layers the approach does not require OS support.

Operating System Integrated Energy Aware Scratchpad Allocation Strategies for Multiprocess Applications (Bologna, Dortmund)

Various scratchpad allocation strategies have been developed in the past. Most of them target the reduction of energy consumption. These approaches share the necessity of having direct access to the scratchpad memory. In earlier embedded systems this was always true, but with the increasing complexity of tasks systems have to perform, an additional operating system layer between the hardware and the application is becoming mandatory. This work presents an approach to integrate a scratchpad memory manager into the operating system. The goal is to minimize energy consumption. In contrast to previous work, compile time knowledge about the application's behavior is taken into account. A set of fast heuristic allocation methods is proposed in this work. An in-depth study and comparison of achieved energy savings and cycle reductions was performed.

Energy Efficient Cooperative Scheduling and Memory Allocation Techniques for Multiprocess Systems (Bologna, Dortmund)

The increasing amount of functionality in contemporary embedded systems implies the usage of complex software where execution of multiple processes is the common case. Usually the processes are interrupted at an arbitrary point in time. In such a scenario the energy savings achieved by utilization of small and therefore fast and energy efficient scratchpad memories could easily be diminished by excessive copy overhead on each context switch. Therefore this work tackles this problem by using compile time knowledge and profiling result to define energy and runtime efficient points in the code, where a context switch could be performed with least overhead. The work presented here applies source-level transformations to the code through insertion of context switch points. Basically it provides cooperative scheduling at source-level under the constraint of a preferred time-slice length, guarantied maximum deviation from this time-slice length and energy efficient placement of context switch points. The source-level compile-time transformations have been developed at the University of Dortmund. A new lightweight scheduling layer has been implemented for the MPARM simulation platform from the University of Bologna. This setup has been used for gathering required profile data and final runtime results.

Worst-case execution-time aware compilation (Dortmund, AbsInt)

The two partners cooperated on establishing a link between the two subdomains of this cluster. The integration of tools from the two domains led to first results, which were published [Fal07, Lok07]. Details are described in the compiler cluster report.

Resource aware design space exploration (DTU, ETH Zürich)

The Technical University of Denmark (DTU) has focussed its activities in resource aware design space exploration on run-time resource optimization. Based on an extension of our ARTS multiprocessor simulation framework which allows for handling dynamic refonfiguration, which accounts for both communication and reconfiguration overhead, DTU have conducted a set of experiments aimed at gaining a better understanding of the dynamic behavior of coprocessor-coupled reconfigurable systems. The first study has been based on an MP3 decoder application and a simple "worst case" resource management algorithm which enforces many run-time reallocations of subsets of the application and, hence, many reconfigurations. The study has focused on coprocessor-coupled architectures where the architecture is partitioned into a homogeneous array of reconfigurable unites (RUs). DTU have studied the impact of different numbers and sizes of RUs, as well as the number of reconfiguration contexts on each RU and the granularity of the RU, i.e. fine or coarse grained, on the run-time behavior of the system. The conclusion from this study [WuMa07] showed that it is possible to

gain performance from such architectures. Based on these experiments, DTU has explored various run-time resource management polices and how they impact the system performance. The results of these experiments [Wu07] have been submitted to the International Conference on Field-Programmable Technology 2007. This work will be continued.

The work on multi-objective design space exploration environment based on the PISA environment for multi-objective optimization from the group of Lothar Thiele, ETH Zurich, was completed with an invited talk at DIPES 2006, October 2006, Braga, Portugal and the publication [Mad06].

Interfaces for real-time components (EPFL Lausanne, ETH Zürich)

According to the workplan, there have been major activities in the area of interfaces for real-time components. EPFL and ETH Zürich have continued working on developing interface formalisms and algorithms for interface compatibility checking for interfaces that expose timing and resource constraints of components. Concretely, the partners hope to understand better the differences and commonalities between their interface formalisms, in order to combine or generalize them. There have been visits from ETH Zurich to EPFL Lausanne and a presentation of the concept of Modular Performance Analysis with Interfaces by Nikolay Stoimenov of ETH Zurich. Particular results of the cooperation are described in the publications listed in section 2.3.2. In partcicular, we have been able to formally describe the interface algebra that has been used at ETH Zurich in terms of the notation introduced by Henzinger and D'Alfaro. Finally, the new concepts could be applied to interface-based rate analysis of embedded systems.

As a result of these discussions, there is a joint participation of both groups in the FP7 project COMBEST (lead by Joseph Sifakis) which shows the achieved degree of integration.

Timing Analysis and Timing Predictability (USaar and AbsInt)

First hard analytical results have been obtained about the predictability of architectural features, in this case cache replacement strategies. These show that the replacement strategy has a strong influence on the precision of any type of cache analysis.

The formal derivation of abstract processor timing models has been mostly implemented. This process starts from a specification of the hardware architecture in VHDL and proceeds by a series of analyses and transformations. Analyses of such models for several kinds of proerties will be possible once formally derived abstract architectural models are available.

Preemptive scheduling of hard real-time tasks requires precise estimations of context-switch costs. These are largely dependent on the cache-refill costs caused by pre-empting tasks. An approach has been developed and implemented that estimates and even minimizes the cache interference of tasks. The latter optimization uses the memory allocation to define the cache mapping.

An integration of AbsInt's aiT timing-analysis tool with the ASCET specification and synthesis tool of ETAS has been realized, and experimental results about the effect have been produced.

Final Results

Fault-tolerant Embedded Systems Design (DTU, Linköping)

As part of the collaboration between Linköping University and DTU and as a continuation of the work in the previous years, an approach to the synthesis of fault-tolerant schedules for embedded applications with soft and hard real-time constraints has been developed. The goal is to guarantee the deadlines for the hard processes even in the case of faults, while

maximizing the overall utility. Time/utility functions are used to capture the utility of soft processes. Process re-execution is employed to recover from multiple faults. A single static schedule computed off-line is not fault tolerant and is pessimistic in terms of utility, while a purely online approach, which computes a new schedule every time a process fails or completes, incurs an unacceptable overhead. Thus, a quasi-static scheduling strategy is used, where a set of schedules is synthesized off-line and, at run time, the scheduler will select the right schedule based on the occurrence of faults and the actual execution times of processes. Moreover, a preemption technique is elaborated as a method to generate flexible schedules that maximize the overall utility for the average case while guarantee timing constraints in the worst case. The scheduling algorithm determines off-line when to preempt and when to resurrect processes. Prof. Paul Pop from DTU has visited Linköping with several occasions during this period. The following joint publications have been generated: [Izo08], [Ele08], [Izo2-08].

Resource analysis using price timed automata (DTU)

DTU has done initial work aimed at using the formalism of price timed automata to model and analyze resources of simple embedded systems. Analysis of resource consumption of embedded systems is a major challenge. DTU has focused on the analysis of timing and memory access costs of those models, using UPPAAL Cora. Although the experiments performed are on small and rather basic models, we have demonstrated [OBH08] that priced timed automata can be used to reason about resource consumption of embedded systems.

Run-time resource management (DTU)

DTU has continued its work on run-time resource optimization. The experiments on dynamic reconfiguration on a coprocessor-coupled architecture, based on an extension of our ARTS multiprocessor simulation framework, have been continued in year 4. Experiments with various run-time resource management polices have shown that it is possible to gain performance from such architectures and have suggested some general guidelines for obtaining efficient run-time resource management [WuHaMa08].

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

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

A master student from Bologna, Paolo Burgio, has spent the period June 2007 – April 2008 at Linköping working, in particular, on controller design. The Symta/P tool for WCET analysis, from Braunschweig, has been used in this project.

Publications in the last period: [And 08], [Ros 07]. Joint publication is under way.

Integration of MPARM and MEMSIM into one simulation platform (Dortmund, Bologna)

MPARM is a multiprocessor simulation platform developed at the University of Bologna. The simulator offers a cycle-true simulation of a ARM7 based multiprocessor SoC. The simulation includes a configurable number of processors and associated memories. Furthermore, additional hardware devices can be simulated. To some extend the memory hierarchy of each processor can be configured. Presence of a caches, as well as size and number of scratchpads can be configured, but there are obvious limitations to the reconfigurability of the whole system. For example, only an one-level cache hierarchy is available. Another restriction forces every processor to implement the same memory hierarchy.

Complementary, at TU Dortmund a versatile memory hierarchy simulator has been developed. It offers the simulation of a free configurable multi-level memory hierarchy. The user provides a memory hierarchy description and an processor-based memory access sequence to MEMSIM, which computes the accumulated timing and energy values for that input. Limitations of that approach are the simulation of a single execution trace (i. e. single processor system) and the missing simulation of the processing unit.

An integration of both simulation platforms has been accomplished in the context of Artist. The resulting MPARM/MEMSIM simulator is capable of simulating a ARM-based multiprocessor SoC with various memory hierarchies. Each processing unit may have a different memory hierarchy and different cycle and energy values. The most challenging task of merge both simulation platforms while ensuring a cycle true execution in the simulation environment has been successfully performed. The resulting simulator offers a shared memory platform with one AMBA bus interconnect. Up to 32 processing units with MEMSIM-defined local memory hierarchies (i. e. local caches, scratchpads etc.) can be attached to this bus. The bus-attached main memory of the original MPARM has been replaced with another MEMSIM component, which allows for defining memory hierarchies instead of plain SRAMs. Since this versatile configurable memory hierarchy is not covered by the access values provided in MPARM, an interface to CACTI has been integrated to automatically generate reliable values for a particular memory hierarchy.

Memory hierarchy design-space-exploration for the MPARM/MEMSIM Simulator (Dortmunt, Bologna)

Based on this new simulation platform, an approach for automatic design space exploration has been developed. The approaches utilizes SPEA2 – an advanced genetic algorithm – for optimizing the memory hierarchy for a particular application. The memory hierarchy can be optimized according to user defined weights in the 3-dimensional space: cycles - energy - die-size. The algorithm produces Pareto-optimal solutions, where from according to the wights the most efficient is selected. During the optimization run thousands of simulation runs with various memory hierarchies may become necessary for the estimation of the fitness of individuals in the genetic pool. For this task the MPARM/MEMSIM simulator has been successfully utilized. The results show that the approach was capable of finding a well suited, efficient and reasonable memory hierarchy for a particular application.

Parallelization of MIMO benchmark using MAPS tools (Aachen, Bologna)

MAPS (MPSoC Application Programming Studio) tools have been developed at RWTH Aachen aiming to help the developers efficiently parallelize the software for embedded parallel architectures. In this activity the MIMO benchmark from Bologna has been evaluated and

parallelized using the MAPS compiler. Unlike a fully autonomous parallel compiler, MAPS provides a set of tools and technologies which helps the parallelization process, for example, advanced source-level profiling and data-flow driven partitioning. A novel concept, called CB (Coupled Block), is proposed to capture the coarse-grained task parallelism on the compiler IR level. TCT (Tightly-Coupled Thread) MPSoC, from Aachen's partner TokyoTech, has been used in the framework (shown in the Fig. 1) as a backend to evaluate the parallelization result. The experiment showed that MAPS could quickly spot the parallelism opportunities in the sequential application and raise the partitioning suggestions to the designer as good starting points. Together with the array analysis capability for the critical loops, the designer could understand the dependency flow easily, which is usually time-consuming in the manual parallelization process. Overall the parallelization productivity is improved. More details of the MAPS tools can be found in [1] and [2].

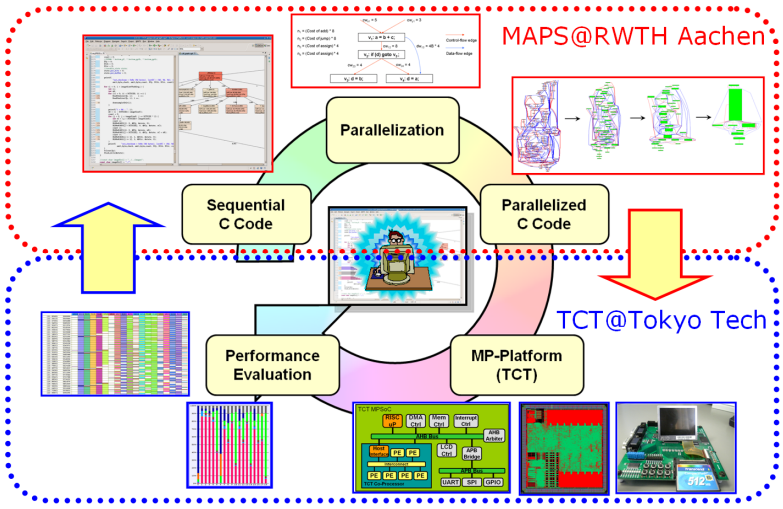

Fig. 1 MAPS-TCT Framework

Resource Aware Design Space Exploration and Mapping (ETHZ, BOLOGNA)

ETHZ has been working together with BOLOGNA in order to establish a joint methodology in order to map algorithms onto multi-processor platforms under resource constraints. In particular, the MPARM environment from BOLOGNA will be used as a simulation backend. The concrete results are as follows:

- Specification of the application using a process network, including an XML based specification syntax.
- Specification of the target architecture, again using a XML description.
- Specification of the mapping while taking different scheduling disciplines into account. A corresponding XML format has been defined.
- Mutli-objective design space exploration using a component-based analysis methodology named MPA-RTC (Modular Performance Analysis – Real-Time Calculus).
- Joined work on the generation of the necessary hardware-dependent software and the compilation towards the target platform has started.

Formal Performance Analysis and Resource-Aware Mapping (ETHZ, EPF Lausanne)

ETHZ has been extending the modular performance analysis method using results from interface theory as provided by EPF Lausanne. This way, it is possible to define resource-aware component interfaces. In addition, ETHZ has been developing a method that allows to map dynamic applications optimally to multi-processor platforms (MPSoC) under static (leakage) and dynamic energy requirements. This method has been intensively discussed with BOLOGNA and comparisons have been made with the methodology based on Benders decomposition (ILP and constraint programming).

Timing Analysis and Timing Predictability (USaar and AbsInt)

The notion of predictability of cache architectures has been clarified. This is the first precise notion of predictability found in the literature. It turns out that the cache-replacement strategy is the decisive characteristic for the predictability of architecture. 4 different replacement strategies were compared, and the LRU strategy was found to be optimal.

The PREDATOR project in the 7th Framework Programme attempts to reconcile performance and predictability. It has identified the PROMPT (PREdictability Of Mulit-processor Timing) design rules for predictable multi-processor design. The first principles are to avoid interference on shared resources in the architecture and to allow the application designer the mapping of applications to target architecture without the introduction of new interferences that were not present in the application.

Parametric timing analysis (USaar and Mälardalen)

Timing analyses require that information such as bounds on the maximum numbers of loop iterations are known statically, i.e., during design time. Parametric timing analysis softens these requirements: it yields symbolic formulas instead of single numeric values representing the upper bound on the task's execution time. So, some input parameters to the program can remain unknown until the final use of the task. The developed analysis determines the parameters of the program, constructs parametric loop bounds, takes processor behaviour into account and attains a formula automatically.

Synergy between Code Synthesis and Timing Analysis (USaar and AbsInt)

One of the problems to be solved synergetically by code synthesis compiling, and timing-analysis is to support mode-specific timing analyses. Many embedded control systems have several operating modes with different timing requirements. Some operating modes are not explicitly specified on the model level or in comments on the C level. Saarland University in cooperation with Bosch (within the PREDATOR project), attempts to semi-automatically identify such operating modes. This would allow timing analysis to implement mode-specific execution-time bounds which can be used to improve schedulability of task sets.

WCET Analysis for Systems with preemptive scheduling (USaar and Absint)

Derivation of timing guarantees was extended in order to cope with the systems with preemptive scheduling. A new method to compute valid upper bounds on a task's worst case execution time (WCET) under preemption was proposed.. This method approximates an optimal memory layout such that the set of possibly evicted cache-entries during preemption is minimized. This set then delivers information to bound the execution time of tasks under preemption in an adopted WCET analysis.

Communication-centric Systems (Cluster Integration)
Led by Rolf Ernst (TU Braunschweig)

Partner teams (leaders): Lothar Thiele – TIK, ETH Zürich (Switzerland), Petru Eles – ESLAB, Linköping University (Sweden), Rolf Ernst – IDA, TU Braunschweig (Germany), Luca Benini – Micrel Lab, University of Bologna (Italy), Jan Madsen – Technical University of Denmark (Denmark)

Affiliated teams (leaders): Fabian Wolf – Volkswagen AG (Germany), Magnus Hellring – Volvo (Sweden), Kai Richter – Symtavision (Germany), Sharon Hu – University of Notre Dame (USA)

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

Work in Year 1

Mixed Performance Analysis in Communication Centric Systems

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

Hybrid Approach for Performance Analysis of Communication Centric Embedded Systems

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

Performance Analysis in the System Design Process

In February 2004, Ernesto Wandeler spent 10 days at the ESI. During this time, an appropriate case-study system was identified and analyzed using a formal performance analysis method developed at TIK, ETH Zürich (Real-Time Calculus). Further, Ernesto Wandeler held 3 talks at ESI, to introduce people at ESI to the performance analysis research at TIK. In April 2005,
Marcel Verhoef spent 5 days at ETH. During this time, a journal paper was written, based on a former conference paper. Further, new potential case-study systems, as well as plans for a performance analysis tool were discussed.

As a first case study, an existing distributed in-car radio navigation system was chosen and was specified in UML. For this case study, Real-Time Calculus was used to evaluate and compare 5 different potential system architectures. Sensitivity analysis was applied to all architectures to identify their robustness and potential bottlenecks. For the architecture that is actually used in the commercial implementation of the case study system, the robustness and the bottlenecks could be identified correctly using formal performance analysis methods.

Optimization and analysis of distributed embedded systems

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

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

Power analysis and optimization

To initiate the joint activity, Bren Mochocki, University of Notre Dame, spent 2 month at TU Braunschweig. During this time, interfaces between the power analysis tools developed at University of Notre Dame and SymTA/S were created. Based on these interfaces SymTA/S could be extended with an analysis technique to determine the power consumption of a given embedded system. Since then the power models and the interfaces were regularly extended and refined.

On-chip interconnections for single-chip execution platforms

The work on communication-centric systems by the group of Prof. Madsen, Technical

University of Denmark (DTU), has focused on on-chip interconnections for single-chip execution platforms. The starting point for this work has been the development of a clock less

NoC architecture (MANGO) and a system-level NoC model based on the multiprocessor simulation environment (ARTS) developed at DTU. The MANGO NoC architecture is based on asynchronous message-passing and provides guaranteed services. Its interface is based on the standard OCP interface protocol which makes the architecture very suited for a modular

SoC design flow. The use of a clockless circuit technique has a number of advantages, among which are; inherent global timing closure, low forward latency in pipelines, and zero dynamic idle power consumption. The ARTS modelling framework, which is developed as part of the

System Modelling Infrastructure activity of ARTIST2, was extended with capabilities to model interconnect structures. At the system-level, the details of the processing elements and the

NoC need to be abstracted in a way that allows for an accurate modelling of the global performance of the system, including the interrelationships among the diverse processors, software processes and physical interfaces and interconnections. To support the designer of single-chip based embedded systems, which includes multi-processor platforms running dedicated RTOS's, with the ability to analyse effects of on-chip interconnect network, the ARTS framework was required to support the analysis of network performance under different traffic and load conditions. This was achieved by extending the model with capabilities for NoC modelling.

Highly scalable communication architectures

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

High level modelling of system interconnects

High level models for system interconnects were at first derived for state-of-the-art busses, validated on a virtual platform and potentials for their deployment were explored. In particular, cooperation with Linköping University paved the way for exploring different high level models for shared communication resources and for exploiting them within theoretical frameworks for efficient allocation, mapping and scheduling of tasks onto MPSoC hardware platforms.

Specifically, we identified two modelling approaches to on-chip communication (additive models and coarse-grain modelling of communication tasks) and addressed the issue of modelling implicit communication in a predictive way (cache misses, semaphore polling). Then, we developed the theoretical framework taking the Benders Decomposition approach: an

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

Hybrid system-level performance analysis approach

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

Work in Year 2

Timing Analysis of the Flexray Protocol

In the second year the Linköping group has continued the work regarding analysis and optimization of distributed embedded real-time systems, with application in automotive electronics. The main goal is to develop models and tools for the analysis and optimization of such communication-intensive systems. Emphasis is placed on the analysis of timing properties, considering the heterogeneous nature of such systems and the particularities of the various communication protocols. In the most recent research the analysis of mixed static/dynamic protocols, such as FlexRay, has been performed [PPE+06].

FlexRay is likely to become a standard for certain automotive applications and the elaboration of the first timing analysis approach for distributed systems built on FlexRay is of importance for our industrial partners. On top of these timing analysis approaches, various system-level optimization tools have been built, performing application mapping, communication synthesis, priority assignment, etc. The Linköping group has closely collaborated with our industrial partners at Volvo as well as with the Braunschweig group. The developed analysis approaches are under integration in the Symta/S environment developed at Braunschweig.

Fault Tolerance

One other issue that has been explored by the Linköping group, in the same context of distributed communication-intensive real-time systems, is that of fault tolerance and, in particular, the issue of transient faults. There are two main aspects of interest here:

(1) Analysis of timing properties in the presence of faults and possible guarantees regarding worst case behaviour

(2) System optimization, such that timing and fault tolerance requirements are satisfied given a certain, limited amount of resources.

An approach for scheduling and worst case analysis with fault tolerance has been developed [IPE+06]. On top of this analysis approach, an optimization technique for task mapping and fault tolerance policy assignment has been elaborated and implemented.

Combination of performance analysis methods: SymTA/S and MPA

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

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

Sureal-Project: Hierarchical Event Models

The main goal of the Sureal Project is to define an integrated development process for distributed embedded real-time Systems, especially regarding real-time aspect in all phases of the development. This includes the integration of different techniques for describing, analyzing and modelling real-time aspects. To be able to use different tools specialized in handling realtime aspects in different phases of the system development interfaces must be defined for them to efficiently work together.

Also the early prediction of the timing behaviour, the sensitivity and optimizing possibilities of the architecture play a very important role in such an integrated development process. The tool SymTA/S is capable of analysing such aspects but the underlying methods still have some limitations regarding specific system setups. Up to date, only task sets, which consist of tasks that are activated according to a standard event model can be analysed appropriately. To lift this limitation, first steps towards exploring hierarchical event models are taken. Future Results will be integrated into SymTA/S to further enhance its applicability.

Power Optimization under Timing Constraints
Based on the power analysis extension to SymTA/S which was realized in cooperation with

Bren Mochocki during the first project year, TU Braunschweig and University of Notre Dame implemented heuristic and stochastic power optimization algorithms using DVS and SVS (Dynamic/Static Voltage Scaling). The presented algorithms are applicable to complex distributed systems with complex timing constraints (maximum jitter, end-to-end deadlines,

etc.), and are capable of determining Pareto-optimal design trade-offs between system power consumption and timing properties.

The heuristic power optimization approach is based on research of TU Braunschweig related to sensitivity analysis [RHE06], whereas the stochastic algorithms utilize the compositional SymTA/S design space exploration framework [HRJ+06].

The results of this activity lead to a joint publication at the International Conference on Compilers, Architectures, and Synthesis for Embedded Systems (CASES) [RHE+06].

Robustness Optimization for Distributed Embedded Systems

Based on the results achieved in the domain of sensitivity analysis [RHE06], TU Braunschweig developed techniques for optimizing the robustness of embedded real-time systems with

respect to variations of system properties like worst-case execution/communication times, bus bandwidth, CPU clock rate, input data rate, etc. Reasons for such variation during the design process or in the field include updates, bug fixes, late feature requests, and product variants.

The developed algorithms consider hard-real time constraints and are capable of optimizing a given system for static and dynamic design robustness. Thereby, the static design robustness optimization approach is applicable to the design scenario where system parameters are fixed early in the design process, whereas dynamic design robustness optimization approach includes possible counteractions to unforeseen system property changes, and is thus applicable to reconfigurable systems.

The results of this research will be published at the International Conference on Hardware - Software Codesign and System Synthesis 2006 (CODES) [HRE06].

Flex Film: High-resolution Real-time Digital Film Applications

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

Simulation-based analysis of SoC interconnection architectures

Industrial MPSoC platforms exhibit increasing communication needs while not yet reverting to revolutionary solutions such as networks-on-chip. The limited scalability of shared busses is being overcome by means of multi-layer communication architectures.

However, the complex interaction among system components and the dependency of macroscopic performance metrics on fine-grain protocol features stress the importance of highly accurate modelling and analysis tools. The work in this area has focused on developing accurate functional model of multi-node on-chip interconnects, as they are currently deployed in high-complexity SoCs today.

Network-on-chip architectures

In the second year the group at the Technical University of Denmark has further developed the

NoC architecture called MANGO (*Message-passing Asynchronous Network-on-Chip providing Guaranteed services over OCP interfaces*). In particular the network core, i.e. the routers and links. MANGO is based on clockless circuit techniques, and thus inherently supports a GALS (*Globally Asynchronous Locally Synchronous*) type design flow. This is an advantage in large scale SoC design, since the distribution of a global clock is becoming increasingly difficult.

MANGO employs virtual channels to provide connection-less best-effort routing as well as connection-oriented virtual circuits, for which service guarantees can be given. The predictability of guaranteed services is a way to promote system-level integrity. The MANGO architecture has been demonstrated through a circuit-level design of a 5x5 router using a 0.13µm CMOS standard cell library from STMicroelectronics. Netlist simulations showed a performance of 650 Mflits/s under typical timing conditions [BS06]. Three patents [BS05] on the MANGO technology have been filed and a startup company, called Teklatech (www.teklatech.com), was formed as a spin-off from this research. Teklatech is developing a

one-step EDA solution to achieving timing closure in large scale, globally synchronous, deep submicron ASIC designs.

Distributed wireless sensor networks

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

Simulation-based analysis of SoC interconnection traffic

In Multi-Processor System-on-Chip (MPSoC) design stages, accurate modeling of IP behavior is crucial to analyze interconnect effectiveness. However, parallel development of components may cause IP core models to be still unavailable when tuning communication performance.

Traditionally, synthetic traffic generators have been used to overcome such an issue. However, target applications increasingly present non-trivial execution flows and synchronization patterns, especially in presence of underlying operating systems and when exploiting interrupt facilities. This property makes it very difficult to generate realistic test traffic. Technical

University of Denmark and University of Bologna have jointly developed a reactive traffic generator device [MAMBS05] capable of correctly replicating complex software behaviours in the MPSoC design phase. The approach has been validated by showing cycle-accurate reproduction of a previously traced application flow. Even when tested under complex synchronization scenarios, including asynchronous interrupts involving OS interaction in a multiprocessor environment, the proposed traffic generator is able to reproduce IP traffic with full capability to express the application flow.

Work in Year 3

Comparison of different performance analysis approaches (ETH Zurich)

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

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

There has been an intensive cooperation between TU Braunschweig and ETH Zurich in terms of coupling their respective frameworks for performance analysis (Symta/S and MPA). As a result, one can now mix different model paradigms in a single analysis run and therefore,

exploit the different modeling scopes and accuracies. The cooperation has led to a joint publication at CODES 2007 [KHE+07].

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

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

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

University of Linköping and DTU have continued the work regarding analysis and optimization of distributed embedded real-time systems, with application in automotive electronics. The main goal is to develop models and tools for the analysis and optimization of such communication-intensive systems. Emphasis is placed on the analysis of timing properties, considering the heterogeneous nature of such systems and the particularities of the various communication protocols.

In the most recent research the Technical University of Denmark (DTU) and Linköping University (LiU) have proposed an approach to the timing analysis of applications communicating over a FlexRay bus [PPE+06], taking into consideration the specific aspects of this protocol, including the dynamic segment. More exactly, they have proposed techniques for determining the timing properties of messages transmitted in the static and the dynamic segments of a FlexRay communication cycle. The analysis techniques for messages are integrated in the context of a holistic schedulability analysis algorithm that computes the worst-case response times of all the tasks and messages in the system.

This was the first step towards enabling the use of this protocol in a systematic way for time critical applications. As a second step, DTU and University of Linköping have proposed an optimization approach for determining a FlexRay bus configuration which is adapted to the particular features of an application and guarantees that all time constraints are satisfied [PPE+07b]. Heuristics for solving the bus access optimization problem with FlexRay have been proposed as part of this research. While the optimization techniques proposed by us can also be applied to other heterogeneous distributed applications, solving the particular problem of analysis and optimization of FlexRay-based systems is, today, of particular importance for the automotive industry.

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

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

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

The Linkoping group has investigated the impact of memory access over shared buses on the global predictability of real-time systems implemented on multiprocessor SoC architectures. We have elaborated an approach [RAE+07] to system level scheduling and bus access which provides predictable implementations in the context of potential bus conflicts. Bus traffic taken into consideration is not only that for inter-task communication but also that caused by regular memory access in the case of cache misses. The basic WCET analysis used as part of the proposed systems is based on the Symta/P tool from Braunschweig.

Network-on-chip architectures (DTU)

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

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

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

Fault tolerance and robustness of autonomous systems (TU Braunschweig)

The AIS project ('Autonomous Integrated Systems') is a research cooperation involving several universities. The aim is to research new system design methodologies which enable MpSoCs to realize autonomous behaviour in case of faults and modifications in the environment. This includes self-healing and self-configuration properties on hardware layer due to specially designed components as well as on software layer. The measures on software layer range from an offline design space exploration at early design stages to the adoption of an autonomous operating system, which supports dynamic reconfiguration mechanism at runtime.

Our main research interest focuses on the determination of the performance of such systems under real-time conditions. Therefore we consider two generally different types of failures. Transient errors have a bounded duration, and they are countered usually with one-time actions. Examples of transient errors are single event upsets or timing failures due to the difficulty of accurate prediction of transistor behaviour. In contrast permanent errors have a potentially unbounded duration, and mostly they affect complete function units or processors. A defect during manufacturing process causing a processor breakdown as well as a permanent performance reduction due to heat fluctuations are exemplary situations for permanent errors.

We implemented algorithms to capture the consequences of this kind of failures and to explore several countermeasures. In this context we developed an approach to analyse the sensitivity of system against performance fluctuations up to complete processor breakdowns. Furthermore possible task migrations have been explored to make the system more robust against errors by remapping functionality within the system in error case. Currently, the system properties during the occurrence of transient errors are investigated, including a derivation of deadline miss probabilities. All analysis algorithms are based on the SymTA/S tool for performance analysis in complex real-time systems. http://www.edacentrum.de/ais

Online performance analysis of distributed embedded systems (TU Braunschweig)

In order to perform online performance analysis in distributed embedded real-time systems, we implemented a framework for distributed analysis of a formal system model [SHE06a, SHE06b]. We based our work on the compositional performance verification approach of Richter et al. as implemented in the tool SymTA/S. A prototype implementation extending the tool has already been tested.

Online Performance control is achieved by using this framework to continuously evaluate the current state of an embedded system. For this, applications entering the system must enter a service contract, stating the computational and communication needs, which are integrated in the current system model. If the resulting system does not violate any constraints, the new application can be accepted into the system. Watchdogs will ensure that no application breaks its service contract. At this time, the concept has been finalized and a prototype has been implemented on a standard PC running parallel processes.

The framework will also be used for optimization of the current system configuration. Heuristic optimization algorithms can perform what-if analyses to continously improve e.g., the robustness of the system. Similar strategies can be applied to deal with faulty hardware, such as degrading processors or if applications break their service contract in order to remain in a feasable system state.

Integrated development process for distributed embedded systems (TU Braunschweig)

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

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

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

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

Applications designed for the MORPHEUS platform may require a massive amount of memory, as well as sufficient bandwidth, to fully demonstrate MORPHEUS's potential as a high-performance reconfigurable architecture. For example, the film grain noise reduction application requires massive amounts of bandwidth due to its real-time requirements. To meet these constraints and to eliminate external memory bottlenecks, a high-bandwidth DDR-SDRAM memory controller has been designed for use with the MORPHEUS platform. Using a

bandwidth-optimizing two-stage buffered access scheduler, bank interleaving and request bundling, the controller achieves read and write throughput levels up to 2 GiB/s.

Real-time Digital Film Noise Reduction (TU Braunschweig)

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

Final Results

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

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

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

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

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

TU Braunschweig has continued the work on modeling of COM-layers in distributed real-time systems with the use of hierarchical event models. The basic research was published in [RE08a], and [RE08b], introducing the hierarchical models. This model allows merging several single event streams to a hierarchical event stream. When the hierarchical event stream is transformed, e.g. due to the fact that it passes different system components, the effects on the previously merged streams can be calculated, which allows the later extraction of the single event streams. These improved models highly improve the accuracy of communication modelling complex distributed real-time systems. The methods have been implemented into the tool SymTA/S.

Online performance analysis of distributed embedded systems (TU Braunschweig)

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

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

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

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

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

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

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

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

As part of its activities in the present project, the Linkoping group has studied the issue of predictable implementation of real-time applications on multiprocessor systems. One of the main difficulties in this context is the problem of memory access due to the shared resources accessed by parallel tasks. This makes the memory access time (in case of cache misses) very difficult to predict in a way that is not overly pessimistic. The Linköping group has elaborated an approach to system level scheduling and bus access which provides predictable implementations. The basic WCET analysis used as part of the proposed systems is based on the Symta/P tool from Braunschweig.

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

Analysis of Distributed Wireless Sensor Networks (ETH Zurich)

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

Generalizing Communication-Centric Performance Analysis (ETH Zurich)

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

FlexRay fault-tolerance aspects (DTU, LiU)

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

NoC arcitectures and programming models (DTU, Bologna)

DTU and Bologna have continued their work on efficient NoC architectures [Mah08]. In particular, DTU has extended their work to include programming models for NoC-based systems, with the aim to address the system designer's point of view. This has been done through a case study of an embedded image processing application [RaStKa08], for which parallelism and scalability has been investigated. The major challenges faced when parallelizing the application were to extract enough parallelism from the application and to reduce load imbalance. The application had limited immediately available parallelism. It was difficult to further extract parallelism since the application had small data sets and

parallelization overhead were relatively high. There was also a fair amount of load imbalance, which was made worse by an non-uniform memory latency. Even so, we have shown that with some tuning relative speedups in excess of 9 on a 16-core system can be reached.

Design for Low Power (Cluster Integration)
Led by Petru Eles (Linköping University)

Partner teams (leaders): Luca Benini – University of Bologna (Italy), Petru Eles – Linköping University (Sweden), Jan Madsen – Technical University of Denmark (Denmark), Lothar Thiele – ETH Zurich (Switzerland).

Affiliated teams (leaders): Roberto Zafalon – STMicroelectronics (Italy), Salvatore Carta – University of Cagliari (Italy).

Overview: The group of Luca Benini at the University of Bologna (UoB) is one of the leading centers in low power design, focusing on system level power management both from the architectural and from the software viewpoint. In this area, the group has produced a large number of contributions on OS-Based-dynamic power management, memory and communication architecture optimization for low power consumption, low power circuit design, battery-driven power management.

One of the baselines contributions of UoB to the project is a complete power modeling infrastructure both for all components of current MPSoC Platforms and for future Network-on-Chip-based platform.

More recently, UoB has focused on power optimization techniques for multi-processor systems on chip, with special emphasis on static and dynamic power management through voltage and frequency setting.

The group of Jan Madsen at the Technical University of Denmark (DTU) aims at low-power techniques for wireless sensor networks, and it brings significant experience on low-power asynchronous circuit design, as well as analytic and stochastic modeling of power consumption and battery usage. They have developed exploration methods for mapping applications onto heterogeneous multiprocessor platforms based on a multi-objective optimization framework from ETH Zurich, where one of the objectives is power consumption.

The group of Petru Eles at Linkoeping University (LIU) has given important contributions on high-level system modelling of both power and reliability, and on optimization techniques for energy efficient mapping of applications on execution platforms. They have developed approaches for energy efficient implementation of both hard and soft real-time systems.

The group of Lothar Thiele at ETH Zurich (ETHZ) has a long standing experience in the area of sensor networks. In particular, a low power platform including hardware, operating system, middleware and various applications has been developed (BTnode). It is used by many research groups worldwide. Together with the experience in real-time systems and scheduling, this is the basis for the joint effort in the design of low power massively distributed systems.

Work in Year 1

University of Bologna has focused on interconnect optimisation techniques for low power. Several schemes have been developed to instantiate application (platform) specific interconnect architectures for minimum energy consumption. An algorithm for automatic instantiation of multi-hop busses which includes topology generation and bus frequency assignment has been developed in collaboration with Penn State University.

A research effort on energy aware mapping of multi-task applications on multi-processor SoC execution platforms has been also started. The approach is based on variable-voltage processors where execution speed and voltage supply can be independently adapted to the processor's workload. The first result of this effort has been a design space exploration technique that automatically finds pareto points in the power vs. throughput design space. The technique has been tested on streaming-like signal processing applications. This work is conducted in cooperation with Linköping University.

Additionally several extensions to the power modelling infrastructure in the MPARM virtual platform simulators have been developed, including the model for variable frequency and variable voltage cores, as well as a prototype model for estimating the power consumption of IOs and external memories (this work has been performed in cooperation with associate partner STMicroelectronics).

Linköping University has developed a technique for static routing on NoC, with guaranteed delays and arrival probabilities in the presence of transient faults. The approach is based on schedulability analysis of tasks and messages with priority based arbitration. For fault-tolerance, a combination of spatial and temporal redundancy is considered. Reduced communication energy is one of the goals. More recently the analysis of the worst-case buffer space needed has been performed. Based on this analysis, it is possible to develop an approach to buffer space minimization in the context described above.

Linköping University's efforts are also aiming at a more accurate modelling of actual communication and memory techniques used in MP SoC. Such an accurate modelling is needed in order for a system level analysis and optimization to produce useful results. Thus, work is concentrating on: (1) Capturing the background communication due to cache misses in system level models. (2) Capturing the bus load due to system-wide synchronization. Once these modelling issues are solved, different optimization techniques can be used for e.g. task mapping and scheduling, as well as voltage selection. Results can be validated using accurate and fast simulation in the environment developed at Bologna. Another issue which is currently addressed is that of efficient optimization techniques based on advanced constraint solving and mathematical programming techniques. This work is performed in cooperation with the group at University of Bologna.

Technical University of Denmark has started the development of a generic sensor network platform (Hogthrob project) which allows to trade-off hardware and software implementations of the various components of the platform. So far, the focus has been on: (1) Processor design: Low-power design techniques have been investigated, included low-power synthesis (e.g., clock-gating), power modes and de-synchronizing in the context of the OpenCores AVR core (using Synopsys) (2) Power modelling: Simulation-based power modelling and estimation techniques have been investigated. This involves analytic and stochastic modelling of batteries and investigation into the macro-modelling of various hardware components.

Based on the prototype sensor network platform developed within Hogthrob, various test bench programs have been run on an AVR core synthesized on the FPGA and a number of physical measurements have been conducted. Finally, DTU has started investigating how the power modelling and the sensor network modelling can be captured within the multiprocessor simulation environment, ARTS, developed at DTU.

Work in Year 2

Power modelling for complex SoC platforms

The activity has focused on extending system-level energy analysis to highly integrated MPSoC platforms with segmented bus architectures, where the efficiency of bridges and protocol/frequency/size converters comes into play to determine the performance of the system interconnect. We leveraged a close cooperation with associate member STMicroelectronics which provided the models, traffic generators, system specifications and performance requirements. Platforms based on the on-chip communication protocols STBus, AMBA AHB, AMBA AXI-have been modelled and simulated at a very high level of accuracy (cycle-accuracy and bus-signal-accuracy), and compared with mixed AHB/AXI platforms

The original MPARM platform allowed the modelling and simulation of single-node communication architectures (as depicted in Fig.1a). The platform was enhanced with the possibility to extend the modelling capability to a multi-layer architecture, as illustrated in Fig.1b. The first scenario corresponds to low-end real-life platforms, where AMBA AHB, AMBA AXI or STBus are the architectures of choice to accommodate on-chip communication. The MPARM platform can also instantiate a NoC as the communication fabric, by wrapping the masters and slaves with the proper network interfaces. In general, all cores can be wrapped with the native bus interface. More complex MPSoC platforms adopt the communication architecture depicted in Fig.1b. It is a hierarchical infrastructure, where communication takes place at a first level of the hierarchy in the local AMBA AHB layers, and at a second level with the system-level slaves. The AMBA Multi-Layer specification introduced the notion of the interconnect matrix first, by envisioning point-arbitration at the destination slaves. This solution is quite interesting, since it allows a larger scalability than single-node solutions. Unfortunately, fabrication problems arise when the number of input layers increases a lot, since the implementation of the interconnect matrix is mostly combinational. This gives rise to clock frequency limitations and to layout unpredictability. As the level of integration of MPSoCs increases, the illustrated structures cannot satisfy communication requirements any more.

A further increase in communication scalability is exposed by segmented architectures, where a number of busses are interconnected with each other by means of bridges. In this case, the congestion on each bus is greatly decreased, thus favouring lower bus access times, but the latency of bus transactions can be seriously increased because of the multiple steps needed to reach a slave located on a different bus. Bridge traversal latency can significantly contribute to overall communication latency. Similarly, the use of bridges raises power concerns. The use of bridges helps to relieve the scalability limitations of traditional communication architectures, however the associated cost consists of the design of a complex IP block (the bridge itself) which is far from trivial and which can significantly affect system performance and energy. Many times, bridges do not perform only protocol conversion, but also size and frequency conversion. In fact, cores with homogeneous characteristics (i.e., clock frequency, data and address bus width) are typically grouped in the same node, therefore each “segment” of the global communication architecture turns out to be a domain with distinctive features. This obviously increases the bridging cost, since up/down size conversions or frequency conversions all take clock cycles to be carried out.

Another issue concerned the porting of traffic generators in order to make the simulation of complex systems in reasonable time possible. Moreover, this allowed overcoming confidentiality problems related to the intellectual property of communicating actors. STMicroelectronics made available its traffic generators for audio and video IP blocks, allowing us to reproduce on the MPSIM environment the traffic patterns of real-life set-top-box platforms with a high level of accuracy.

Another effect of the joint work on traffic generators between Technical university of Denmark and University of Bologna was the development of the necessary infra-structure to co-simulate modules of the abstract system-level MPSoC ARTS frameworks (DTU) with modules available in the cycle-true MPARM framework. The motivation of the work is to investigate MPSoC instances at mixed-levels of abstraction. A simple system where two ARTS IP cores were connected through a MPARM AMBA-AHB bus was successfully implemented and co-simulated.

Finally, a significant modelling effort was required also for the memory controller. In fact, MPSIM has traditionally simulated MPSoC systems with on-chip memories only; therefore we needed to model real-life memory controllers for I/O. We got the LMI specification from STMicroelectronics, and developed a SystemC model which was accurately (cycle-by-cycle) validated against the behaviour of the real LMI. Such powerful model allows us to interface our MPSoC with SDR and DDR SDRAMs, and more interestingly to model I/O access latency of real systems. Finally, we retain the capability to model an on-chip shared memory in place of the off-chip SDRAM, thus being able to differentiate system performance and power in presence of a slow off-chip memory vs. a fast on-chip memory. Optimizations for access to the off-chip memory can also be analyzed with this platform.

Power optimization via system-level resource allocation and scheduling

In this activity, the focus is on addressing resource allocation problems in Multi-Processor Systems-on-Chip (MPSoCs). An important instance of this problem is when have to allocate and schedule a given task graph (representing a functional abstraction of a multi-task application) on a target multi core platform while choosing the frequency (and voltage) at which each task will be executed. Since hardware platforms and applications are extremely complex, it becomes thus important not only to measure the optimizer efficiency as done in general in the optimization area, but also to verify if the optimization model is accurate through a validation step performed via simulation on a virtual platform.

Allocation, scheduling and discrete voltage selection problem for variable voltage/ frequency MPSoCs, minimizing the system energy dissipation and the overhead for frequency switching, are clearly NP-hard problems. Only incomplete approaches have been proposed to solve these problems in the system design community. In this activity we have investigated a hybrid methodology based both on Constraint Programming (CP) and Integer Programming (IP) that splits the overall problem in two subproblems, the first being the allocation of tasks to processors and frequencies to tasks and the second being the scheduling. Our methodology derives static allocation, scheduling and frequency setting; therefore it targets applications with design-time predictable behaviour.

In order to solve the problem to optimality without incurring accuracy limitations, we applied the concept behind the *logic-based Benders decomposition technique* to this new application problem. Bender decomposition can be summarized as follows. A complex optimization problem is decomposed in two parts: the first, called Master Problem, is the allocation of processors and frequencies to tasks and the second, called Subproblem, is the scheduling of tasks given the static allocation and frequency assignments provided by the master. The master problem is tackled by an Integer Programming solver while the subproblem through a Constraint Programming solver. The two solvers interact via generation of no-goods (constraints on acceptable solutions for the CP solver) and cutting planes (constraints on acceptable values of the integer variables for the IP solver) generation. The solution of the master is passed to the subproblem in an iterative procedure that is proved to converge to the optimal solution. This work has been performed in cooperation by the University of Bologna and Lonköping University.

The methodology has been tested on a variety of realistic instances. In addition, we test the accuracy of the solutions provided by the optimizer simulating them on an MPSoC virtual platform. In particular, we have used two demonstrators (GSM and JPEG) to prove the applicability of the developed methodology to real-life embedded applications scenarios.

In a parallel, but strongly related activity, we have also addressed the specific problems of soft real-time systems. In this case, certain tasks are allowed to miss their deadlines. This however, negatively affects the delivered QoS. The goal is to maximize the QoS with a limited energy budget or to achieve a certain level of QoS with as low energy consumption as possible. Linköping University has developed heuristics which determine the system schedule and voltage levels of tasks in such a system.

Finally, DTU has experimented with the use of meta-heuristics to solve the mapping a set of task graphs onto a heterogeneous multiprocessor platform. The objective is to meet all real-time deadlines subject to minimizing system cost and power consumption, while staying within bounds on local memory sizes and interface buffer sizes. Our approach allows for mapping onto a fixed platform or onto a flexible platform where architectural changes are explored during the mapping. The approach uses multi-objective evolutionary algorithms and is based on the PISA framework for multi-objective optimization developed at ETH Zurich. We demonstrate the approach through an exploration of a smart phone, where five task graphs with a total of 530 tasks after hyper period extension are mapped onto a multiprocessor platform. The results show four non-inferior solutions out of 10.000 explored solutions, which tradeoffs the various objectives.

Scheduling based energy optimization for energy-scavenging wireless sensor networks

Wireless sensor networks – consisting of numerous tiny sensors that are unobtrusively embedded in their environment – have been the subject of intensive research. As for many other battery-operated embedded systems, a sensor's operating time is a crucial design parameter. As electronic systems continue to shrink, however, less energy is storable on-board. Research continues to develop higher energy-density batteries and supercapacitors, but the amount of energy available still severely limits the system's lifespan. As a result, size and weight of most existing sensor nodes are largely dominated by their batteries.

On the other hand, one of the main advantages of wireless sensor networks is their independence of pre-established infrastructure. That is, in most common scenarios, recharging or replacing nodes' batteries is not practical due to (a) inaccessibility and/or (b) sheer number of the sensor nodes. In order for sensor networks to become a ubiquitous part of our environment, alternative power sources should be employed. Therefore, environmental energy harvesting is deemed a promising approach: If nodes are equipped with energy transducers like e.g. solar cells, the generated energy may increase the autonomy of the nodes significantly. Several technologies have been discussed how, e.g., solar, thermal, kinetic or vibrational energy may be extracted from a node's physical environment. Moreover, several prototypes have been presented which demonstrate both feasibility and usefulness of sensors nodes which are powered by solar or vibrational energy.

The focus of this activity is on sensor nodes with energy-scavenging features. In general our results apply for all kind of energy harvesting systems which must schedule processes under deadline constraints. For these systems, new scheduling disciplines must be tailored to the energy-driven nature of the problem. This insight originates from the fact, that energy – contrary to the computation resource "time" – is storable. As a consequence, every time we withdraw energy from the battery to execute a task, we change the state of our scheduling system. That is, after having scheduled a first task the next task will encounter a lower energy level in the system which in turn will affect its own execution. This is not the case in conventional real-time scheduling where time just elapses either used or unused.

The main developments obtained in this activity can be summarized as follows

(a) We studied an energy-driven scheduling scenario for a system whose energy storage is recharged by an environmental source. For this scenario, we developed an optimal online algorithm that dynamically assigns power to arriving tasks. These algorithms are "energy-clairvoyant", i.e., scheduling decisions are driven by the knowledge of the future incoming energy.

(b) We developed an admittance test that decides, whether a set of tasks can be scheduled with the energy produced by the harvesting unit, taking into account both energy and time constraints. For this purpose, we introduced the concept of energy variability characterization curves (EVCC).

(c) In addition, a comparison to earliest-deadline first (EDF) by means of simulation, demonstrated that significant capacity savings can be achieved by our approach, when compared to the classical EDF algorithm.

Work in Year 3

Power optimization and analysis for Nanometer technologies (UoB)

Leakage power is a major concern in sub-90nm CMOS technology and numerous design techniques have been proposed to reduce standby leakage in digital circuits. Out of this rich set of solutions, power-gating or sleep transistor insertion has proven to be a very effective approach to reduce standby leakage, while keeping high speed in the active mode. Another interesting solution is based on the usage of multi-threshold libraries and multiple power supplies. While the first approach is considered to be more effective, it has significant impact on design automation flows. The second approaches is less disruptive in terms of design automation, but it often leads reduced savings because it is applied very late in the design process (i.e. during low-level logic synthesis).

The work performed by UoB aimed at addressing the above mentioned limitations of power-gating and multi-threshold/voltage synthesis. The work has been performed in cooperation with associate partner STMicrolectronics, Bulldast and Politecnico di Torino, in close cooperation with the 6th FP integrated project ICT-CLEAN. The contributions have been in two main directions:

a. Developing design automation techniques for automatic power gating insertion. The following issues have been tackled: (i) clustering of non-timing critical cells in a physical design aware fashion to enable partial power gating with low area overhead. The key idea is to prevent power gating of timing critical cells that would decrease the speed of the circuit. (ii) sleep transistor sizing and insertion. The contribution is to automatically size sleep transistor cells for a specified speed within area constraints. Results on 65nm libraries have demonstrated that significant power savings can be achieved even for circuits that are not amenable to full-power gating because of tight timing constraints (up to 80% of the full power savings potential without any speed degradation has been obtained).

b. Development of high-level (register-transfer level) strategies for automatically exploring the speed vs. power trade-off for arithmetic units in datapath circuits. This work enables a leakage-power aware arithmetic unit instantiation strategy during the early phases of the synthesis process. Given a timing constraint and an average activity information, the new algorithm is capable of selecting the most energy efficient functional unit (e.g. carry lookahead vs. carry propagate adder) in an automatic fashion. Experiments have demonstrated that significant improvements can be achieved in energy efficiency compared with the traditional strategy that optimizes for leakage power only during the

last steps of logic synthesis. In average 20% extra leakage savings have been obtained in a 65nm library with respect to standard logic synthesis techniques.

These work represent the first steps in developing new methods and tools for designing embedded systems in leakage-dominated technologies, a topic that will be of growing interest for the embedded system design community. This work has particular relevance for the ARTIST2 activity **Power optimization via system-level resource allocation and scheduling,** as the high-level allocation and scheduling decisions will be increasingly influenced by leakage power, and by the availability of low-leakage shutdown states.

The main WEB reference to the work performed in this area is the web site of the CLEAN project: http://clean.offis.de The following papers detail the technical achievements obtained with the participation of UoB.

Power optimization via system-level resource allocation and scheduling (LiU, UoB)

Linkoeping and Bologna have continued to the cooperation on this topic, focusing on energy optimization, and more specifically on the optimization of energy-efficient time constrained multiprocessor systems. We tackled optimization problems in system level design for variable voltage/frequency MPSoCs, aiming at minimizing power consumption.

The major challenge that we faced was to develop a cooperative solving framework where the allocation, the scheduling and the discrete voltage/frequency selection problem models can be suitably accommodated and solving algorithms integrated.

By leveraging the principle of logic-based Benders Decomposition we created with an iterative two-step mapping framework that is proved to converge to the optimal solution. Our methodology derives static allocation, scheduling and frequency settings, therefore targets applications with design-time predictable behaviour. Signal processing and multimedia applications employing pipelining as workload allocation policy are the most common example of such applications. Therefore our optimizer was tuned for energy-efficient mapping of pipelined task graphs on MPSoCs. Finally, we used two demonstrators to prove the applicability of the developed methodology to real-life scenarios.

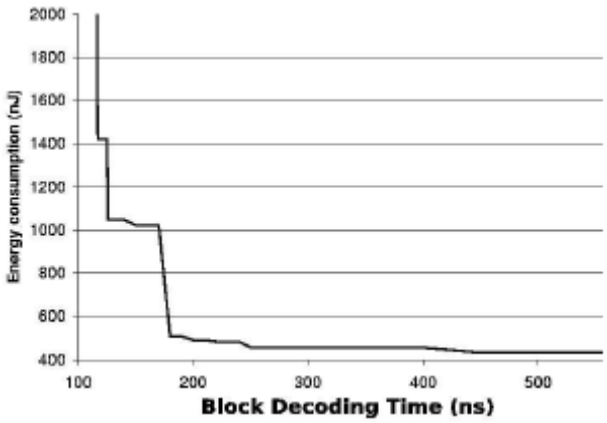

The above figure shows the Pareto-optimal set of solutions, spanning the trade-off between decoding time and Energy consumption for a GSM decoding application running on a multi-core ARM platform. The solution points in the graph where computed by running the optimizer with different block decoding time constraints. As it can bee seen, there is a wide range of feasible operating points that can be automatically computed by the optimizer. It is important to emphasise that both energy and block decoding times have been validate with cycle accurate simulation and power estimation on a virtual platform. Hence, the accuracy of the optimization model has been carefully assessed.

Adaptive Power Management for Energy Harvesting Sensor Nodes (UoB and ETHZ)

Energy is a primary constraint in the design of sensor networks. This fundamental energy constraint further limits everything from data sensing rates and link bandwidth, to node size and weight. For example, it has become evident that energy storage devices largely dominate the form factor of existing sensor nodes and prevent their miniaturization. Moreover, the limitation of the energy supply has constantly impeded the progress of sensor networks towards large scales and true autonomous operation. For instance, long-term monitoring of the environment is one of the sensor network visions. For this application, finite energy reservoirs like batteries render the deployment and maintenance of large scale sensor networks extremely cumbersome.

The focus of this activity is on sensor nodes with energy-scavenging features. Several technologies have been proposed in the past years how, e.g., solar, thermal, kinetic or vibrational energy may be extracted from a node‘s physical environment. Solar energy is certainly one of the most obvious and promising energy sources. Clearly, one may just use solar energy to recharge a primary energy source, like e.g. a battery. Like that, the point in time when the system runs out of energy is simply postponed. If, however, one strives for perpetual operation, common power management techniques have to be reconceived. In addition to perform classical power saving techniques, the sensor node has to adapt to the stochastic nature of the solar energy. Goal of this adaptation is to maximize the utility of the application in a long-term perspective. The resulting mode of operation is sometimes also called *energy neutral* operation: The performance of the application is not predetermined a priori, but adjusted in a best effort manner during runtime and ultimately dictated by the power source.

The main developments achieved in this activity can be summarized as follows:

(a) We present feedback controllers which adapt task rates such that maximal utility is obtained while respecting the time-varying amount of harvested energy. Instead of solving the optimization problem on-line which may be prohibitively complex in terms of running time and power consumption, we propose the use of multiparametric, optimal on-line controllers.

(b) We present a hierarchical control design which overcomes several drawbacks of previously proposed designs: First, the computation overhead and storage demand of the online controller is reduced significantly. In a case study, the achieved reductions compared to a single controller amount 91% and 83%, respectively. Second, by designing the controller for worst-case situations, depletion of the energy storage is avoided and robustness of the overall system is increased.

(c) For some applications of practical concern, the optimal multiparametric solutions may grow to complex for constraint systems like sensor nodes. For those applications, we propose a novel algorithm for approximate multiparametric linear programming. We demonstrate, that the online complexity of the generated control laws is highly reduced compared to an optimal solution in terms of computation overhead and storage demand. For an example investigated, we found improvements of 98% and 92%, respectively. Furthermore, simulations show that the the performance of the found control laws is comparable with the one achieved by complex, optimal control laws.

(d) We propose a practical technique for the efficient implementation of the controller and demonstrate the practical relevance of our approach by measurements of the controller running on a real sensor node.

(e) Beside this theoretical framework, a novel analog circuit for solar energy harvesting has been presented. Here, a key design challenge is how to optimize the efficiency of the solar energy collection under non stationary light conditions. The proposed scavenger exploits miniaturized photovoltaic modules to perform automatic maximum power point tracking.

Energy-aware routing in Wireless Sensor Networks (DTU)

The Technical University of Denmark (DTU) have studied low power aspects of wireless sensor networks, in particular for networks where the nodes are able to harvest energy from the environment and hence being able to recharge their batteries. DTU have developed an adaptable energy-aware routing algorithm based on directed diffusion which is able to avoid routing through nodes with low energy. This strategy allows low energy nodes to recover through energy harvesting. When such a node reaches a certain energy level, it may again be considered for routing data. The additional energy usage introduced by the more advanced routing algorithm is spent for a smart distribution of traffic which takes into account dynamic changes of node status and energy harvesting for the benefit of the network performance and lifetime. The routing algorithm has been implemented and simulated in a setup which uses a solar panel as the energy harvesting unit, and simulates the environment by varying the amount of energy harvested over day and night as well as by the introduction of additional factors like natural shadows which are influencing the amount of energy being harvested. Such kind of fluctuations are noticed by the algorithm and used to arrange traffic and energy consumption to avoid node areas of lower energy.

Communication between nodes in wireless sensor networks is susceptible to transmission errors caused by low signal strength or interference. These errors manifest themselves as lost or corrupt packets. This often leads to retransmission, which in turn results in increased power consumption reducing node and network lifetime. DTU have implemented and evaluated a convolution code FEC with Viterbi decoding on Mica2 nodes to explore the possibility of extending the lifetime of a degrading wireless sensor network. Results from these experiments [DOM06] suggest that the approach could be used in a wireless sensor network when increasing distance and channel noise due to node dropout, degrade the network.

Energy optimization of fault-tolerant embedded systems (LiU, DTU)

Addressing simultaneously energy and reliability is especially challenging because lowering the voltage to reduce energy consumption has been shown to exponentially increase the number of transient faults. The main reason for such an increase is that, for lower voltages, even very low energy particles are likely to create a critical charge that leads to a transient fault. However, this aspect has received very limited attention. Moreover, time-redundancy based fault-tolerance techniques, such as re-execution, and voltage scaling-based low-power techniques are both relying on the use of processor idle-time. In addition, such competing requirements have to be met within a given development and manufacturing cost and time-frame. Therefore, the task of designing such embedded systems is becoming not only increasingly important, but also increasingly difficult.

Linkoeping University (LiU) has proposed a technique to combine re-execution and active replication in an optimized implementation that leads to a schedulable fault-tolerant application without increasing the amount of employed resources. They have also addressed transparency/performance trade-offs during the synthesis of fault-tolerant schedules. Starting from their research, the Technical University of Denmark (DTU) and LiU have jointly considered a very different trade-off, namely, energy versus reliability.

They have proposed an approach to the scheduling and voltage scaling of embedded real-time applications that decides the voltage levels and start times of processes and the transmission times of messages, such that the transient faults are tolerated, the timing constraints of the application are satisfied and the energy consumption in the no-fault scenario is minimized [Pop07a, Pou07b]. This research has considered heterogeneous distributed time-triggered systems, where both processes and messages are statically scheduled. The transient faults are tolerated through process re-execution by switching to pre-determined contingency schedules.

Final Results

Modelling and Optimization of miniaturized Solar Energy Harvester Systems (UoB, ETHZ)

The limited battery life-time of modern embedded systems and mobile devices necessitates frequent battery recharging or replacement therefore the interest in embedded portable systems and wireless sensor networks (WSN) that scavenge energy from the environment has been increasing over the last years. In particular photovoltaic energy harvesting techniques could help to the progress in low-power design and to reduce the size of the battery package in distributed embedded systems. Nowadays small solar panels suffice to ensure continued operation. In Year 3 we proposed a scavenger prototype which exploits miniaturized photovoltaic modules to perform automatic maximum power point tracking. We verified that energy consumption and efficiency of the MPP tracker are very important design criteria in energy scavengers for sensor nodes. Thus, in Year 4, we focused on the optimization of two important metrics maximizing becomes fundamental at low light irradiance:

- Maximization of the energy harvesting efficiency;
- Minimization of the energy used for ineffective operations (e.g. set of operations which cannot be finished for lack of energy) or for too expensive power conditioning processes (e.g. DC/DC components).

In order to fulfil both requirements, the activity, oriented to modelling and optimization of photovoltaic harvesting architectures, can be summarized in the following achievements:

(a) The amount of energy that can be harvested depends on various factors such as the voltage level of the storage device and the incident light intensity. Therefore we investigated the characteristics of supercapacitors as energy storage devices for portable and wearable systems. We find that the voltage across the supercapacitor may be an unreliable indicator of the stored energy if previous charging/discharging cycles or precharging are not taken in account. This result impose a revision of common assumptions of the classical formula $E_{CAP} = 1/2CV^2$ generally used to determine and estimate the energy stored in a scavenging system.

(b) When the size of the photovoltaic module is very small, optimizing the efficiency of energy collection and tracking the Maximum Power Point (MPP) becomes hard and less effective. In fact, Maximum Power Point Tracking technique MPPT is practicable only if the power consumed by the tracking circuit is substantially lower than the amount of output power gained from PV modules. We tackled the challenge of powering a sensor node with miniaturized photovoltaic modules of some mm^2 proposing a new inductor-less architecture for the harvesting process suitable for on-chip integration. The designed scavenger focused only on increasing and optimizing the harvesting efficiency under fixed low solar intensity

(c) We provide guidelines for optimizing solar harvester design and an iterative optimization methodology that allows to improve significantly the whole harvester efficiency. It is worth noting that this methodology is not limited to self powered WSN nodes, but it can be used also to optimize the design of harvesting ICs.

Concerning software to optimize the utility of sensor applications, the methods from the previous year have been further improved. The work addresses both robustness as well as low complexity design for solar powered sensors. Our efforts resulted in two new publications.

A hierarchical control solution has been design which overcomes several drawbacks of previously proposed approaches: First, the computation overhead and storage demand of the on-line controller is reduced significantly. In a case study, the achieved reductions amount 91% and 83%, respectively. Second, by designing the upper control layer for worst-case situations, depletion of the energy storage is avoided and robustness of the overall system is increased. Specifically, a new worst-case energy prediction algorithm accounts for the fact that the short-term energy prediction can be considerably lower than average expectations. And third, the

proposed control formulation renders manual parameter tuning of finite horizon controllers unnecessary, leading to an automatical stabilization of the system. Both simulation results as well as detailed analysis of the implementation overhead on a real sensor node are provided.

A novel algorithm for approximate multiparametric linear programming has also been proposed. The online complexity of the generated control laws is highly reduced compared to an optimal solution in terms of computation overhead and storage demand. For an example application of practical relevance, we found improvements of 95% and 92%, respectively. For many applications, the optimal multiparametric solutions may grow to complex for constraint systems like sensor nodes. Moreover, with increasing complexity, well-established solvers come to their limits and fail to find the optimal solution. For all these applications, the presented algorithm may find useful approximations which exhibit performance metrics comparable to the optimal solution.

Scheduling Based Energy Optimization for Energy-scavenging Wireless Sensor Networks (DTU, UoB)

DTU and Bologna have started a collaboration on energy-scavenging wireless sensor networks. A key goal in wireless sensor network is long node life. Renewable energy sources can potentially lead to perpetual operation, but not without careful management of the energy both in the individual node and in the whole network. DTU has developed a simulator for wireless sensor networks, which is capable of capturing nodes with energy harvesting. DTU has developed a dynamic energy-aware routing protocol, which is able to route traffic within the network such that nodes with limited energy are avoided until they have harvested enough energy. Bologna has together with ETHZ developed a lazy-scheduling algorithm to be used on sensor nodes with energy harvesting. Bologna and DTU have made the first attempts towards integrating lazy-scheduling with energy-aware routing within the simulator.

Temperature Aware System-level Power Optimization (Linköping, UoB)

Linköping and UoB have continued the cooperation on system-level design issues, focusing on energy optimization, and more specifically on the optimization of energy-efficient time constrained multiprocessor systems. In particular, the problem of thermal aware energy optimisation has been tackled.

The major challenge was to integrate temperature modelling into the framework of energy efficient system level scheduling and voltage selection.

High power densities in current SoCs result in both huge energy consumption and increased chip temperature. We have elaborated a temperature-aware dynamic voltage selection technique for energy minimization and performed a thorough analysis of the parameters that influence the potential gains that can be expected from such a technique, compared to a voltage selection approach that ignores temperature. We have also made a study regarding the relevance of taking into consideration transient temperature effects at optimization, the impact of the percentage of leakage power relative to the total power consumed and of the degree to which leakage depends on temperature. Moreover, we have also proposed a temperature-aware task mapping technique for energy optimization in systems with dynamic voltage selection capability.

Over the four years period of the whole project, the cooperation of the involved groups has led to significant results in some of the most relevant issues regarding modern, energy critical embedded systems. They span from issues related to power models for nanometer technology circuits to the ultra low-power nodes for wireless sensor networks:

- *Power optimization and analysis for nanometer technologies:* As result of this work design automation techniques for automatic power gating insertion have been

developed. Results on 65nm libraries have demonstrated that significant power savings can be achieved even for circuits that are not amenable to full-power gating because of tight timing constraints. On top of these circuit-level techniques, high-level (register-transfer level) strategies for automatically exploring the speed vs. power trade-off for arithmetic units in datapath circuits have been elaborated. This work enables a leakage-power aware arithmetic unit instantiation strategy during the early phases of the synthesis process.

- *System level performance and power analysis for multiprocessor systems on chip:* As part of this project, the MPARM platform for SoC simulation and analysis has been extended and has been established as a common working platform for all participating groups. The original MPARM platform allowed the modelling and simulation of single-node communication architectures. The platform was enhanced with the possibility to extend the modelling capability to a multi-layer architecture. Additionally several extensions to the power modelling infrastructure in the MPARM virtual platform simulators have been developed, including the model for variable frequency and variable voltage cores, as well as a prototype model for estimating the power consumption of OS and external memories.
- *System-level energy optimisation of real-time systems:* A research effort on energy aware mapping of multi-task applications on multi-processor SoC execution platforms has been undertaken. The approach is based on variable-voltage processors where execution speed and voltage supply can be independently adapted to the processor's workload. In this activity, the focus was on addressing resource allocation problems in Multi-Processor Systems-on-Chip (MPSoCs). An important instance of this problem is the allocation and scheduling of a given task graph (representing a functional abstraction of a multi-task application) on a target multi core platform while choosing the frequency (and voltage) at which each task will be executed.
- *Sensor networks and energy harvesting sensor nodes:* A generic sensor network platform has been developed which allows to trade-off hardware and software implementations of the various components of the platform. The power aspects of such a network have been studied, in particular for networks where the nodes are able to harvest energy from the environment and hence being able to recharge their batteries. For such systems, new scheduling disciplines must be tailored to the energy-driven nature of the problem. A scavenger prototype was proposed which exploits miniaturized photovoltaic modules to perform automatic maximum power point tracking. Techniques were elaborated for (1) the maximization of the energy harvesting efficiency and (2) the minimization of the energy used for ineffective operations.

2.2.5 Control for Embedded Systems Cluster

This cluster is composed of the following activities:

Design Tools for Embedded Control (Platform)
Led by Martin Törngren (KTH)

Partner teams (leaders): Martin Törngren – KTH (Sweden), DeJiu Chen – KTH (Sweden), Karl-Erik Årzen – LTH (Sweden), Anton Cervin – LTH (Sweden), Zdenek Hanzalek – CTU (Czech Republic), Pedro Albertos – UPVLC (Spain).

Affiliated teams (leaders): Henrik Lönn. Volvo Technology Corporation (Sweden), Jonas Edén, Scania Corporation (Sweden), Diana Malvius, Syntell corporation and KTH (Sweden), Jakob Axelsson - Volvo Car Corporation (Sweden), Ulrich Freund – ETAS Corporation (Germany), Joachim Stroop – dSPACE (Germany), Rolf Johansson – Volcano/Mentor Graphics (Sweden and Hungary), Yiannis Papadopolous – Univ. of Hull (UK), Mikael Strömberg – Systemite Corporation (Sweden), Vladimir Havlena - Honeywell Prague Labs (Cech republic), Yves Sorel – Inria (France), Daniel Simon – INRIA (France), Christoph Kirsch, University Salzburg (Switzerland).

Overview: Several tools have already been developed separately by the individual teams, and are briefly described in the following paragraphs. A national Swedish research programme, FLEXCON (http://www.control.lth.se/FLEXCON/) – which ended 2005, included objectives for integrating these tools, and this JPIA builds on this effort.

Two Matlab-based toolboxes, Jitterbug and TrueTime, for analysis and simulation of real-time control systems have recently been developed at Lund University. The tools can be used at early design stages to determine how sensitive controllers are to scheduling-induced delays and jitter. They can also be used at the implementation stage for trade-off analysis between the tasks. Furthermore, TrueTime can be used as an experimental platform for research on flexible scheduling.

At KTH, the AIDA toolset has been developed for design of networked embedded control systems. The toolset is based on a modelling framework allowing functional requirements and various implementation abstractions to be represented. AIDA supports end-to-end timing behaviour and facilities for fault injection and robustness experiments. Based on experiences with the AIDA toolset, further work has concentrated on developing a new model and tool integration platform.

At CTU, the Torsche (Time Optimisation of Resources, SCHEduling) MATLAB-based toolbox is being developed with support for scheduling algorithms that can be used for applications such as high level synthesis of parallel algorithms and optimized production of manufacturing lines.

UPVLC has developed several co-design tools to facilitate the embedded control system development. These tools include the schedulability analysis of the system with a partitioned system in order to reduce the jitter, optional activities analysis, dynamic changes of controllers and embedded control system generation. RT-LEAST is a tool to deploy minimal embedded control system for RT-Linux.

Work in Year 1

The following is an extract from the 1st year's deliverable, describing accomplishments the first 12 months.

Work achieved in the first 6 months:

- Development of the TrueTime tool (wireless network blocks, battery-powered devices, ocal clocks with drift and offset) – LUND
- Development of a new tool for model integration and management (Paper to appear in the 31st EUROMICRO conference, 2005, by Jad El-khoury, Ola Redell and Martin Törngren) – KTH
- Started the work on a survey on tools for modelling and design of real-time control systems
- Further developments of the TORSCHE (Time Optimisation of Resources, SCHEduling) MATLAB-based toolbox – CTU. Using the toolbox, one can easily and quickly obtain an optimal code of computing intensive applications running on specific hardware architectures like FPGAs with special purpose macros. The tool can also be used to investigate application performance prior to its implementation and to use these values (e.g. the shortest achievable sampling period of the filter implemented on given set of processors) in the control system design process performed in Matlab/Simulink.
- Further development of the tools from UPVLC - UPVLC

Work achieved in months 6-12:

- Completed the survey on tools for modelling and design of real-time control systems. Existing tools have been categorized. In doing so discussions have taken place with the HRT cluster as well as the Hycon NoE to provide feedback on the types of tools included. Continued development of the TrueTime tool (wireless network blocks, battery-powered devices, local clocks with drift and offset) – LUND
- Development of course and training material for TrueTime – KTH
- Tutorial on TrueTime given at IFAC World Congress, Prague, July 3
- Continued development of a new tool for model integration and management (Paper to appear in the 31st EUROMICRO conference, 2005, by Jad El-khoury, Ola Redell and Martin Törngren) – KTH
- Continued developments of the TORSCHE (Time Optimisation of Resources, SCHEduling) MATLAB-based toolbox – CTU
- Continued development of the tools from UPVLC – UPVLC
- KTH has initated a state of the art survey on approaches for model/tool integration and model management

Work in Year 2

Technical Achievements, Outcomes and Difficulties encountered are described in the following – extracted from the year 2 deliverable. Note that the cited references refer to those in the year 2 deliverable.

Achievement: Dissemination of results on design tools to the scientific community

As part of the dissemination of cluster results in this area, we have organized the following events:

- A graduate school on embedded control systems (Prague, April 3-7, 2006)

http://www.artist-embedded.org/FP6/ARTIST2Events/Events/EmbeddedControl/

- A cluster session on Tools for Co-Design of Control Systems and Their Real-Time Implementation at the IEEE International Symposium on Computer-Aided Control Systems

Design (CACSD), Thursday October 5, 2006 (note that the planning for the event took place during the 2nd year of ARTIST2 whereas the event was carried out during the 3rd year).

http://www.elet.polimi.it/conferences/cca06/CACSD_home.htm

The session on “Tools for Co-design of Control Systems and their Real-time Implementation” was prepared by Zdenek Hanzalek, Martin Törngren and Karl-Erik Årzén. The session will be held at the IEEE Conference on Computer Aided Control System Design (CACSD) in Munich, October, 2006.

This session sets the context of embedded control systems development describing what is achievable with current generation tools. The aim of this session is to:

- give overall characteristics of the area
- identify and summarize important co-design tools available
- characterize the state of practice for both industrial and academic tools
- show illustrative case studies
- provoke discussion on integration of these tools.

The session consists of one survey presentations (Tools supporting the co-design of control systems and their real-time implementation; current status and future directions) plus five presentations oriented towards specific tools and principles (Model based integration from the Royal Institute of Technology, Jitterbug and TrueTime from Lund University, Sweden, TORSCHE from the Czech Technical University in Prague, the schedulability issues from Valencia, the SAE Architecture Analysis & Design Language from Carnegie Mellon Software Engineering Institute, US Army/AMCOM and Honeywell Labs).

As part of an effort to summarize achievements in the Swedish research program on embedded real-time systems – ARTES – a chapter was written jointly by KTH and LTH describing the co-design tools that were partly developed by funding from ARTES [9]. See http://www.artes.uu.se/bok/ for more information about the book.

The work has also been promoted and disseminated through a number of invited talks described in section 2.2.3, in some cases coinciding with invited papers [2, 3].

Output from Achievement: Dissemination

- On-line documentation/presentations including overviews of co-design tools – see

http://www.artist-embedded.org/FP6/ARTIST2Events/Events/EmbeddedControl/links

- The papers produced for the CACSD session – see links above and references to individual papers [4, 6, 7, 8].

Difficulties with Achievement: Dissemination

No difficulties encountered.

Achievement: Interactions with other ARTIST2 clusters, and a characterization of model and tool integration efforts

In order to stimulate interactions swith the other clusters, we issued our tool survey for review to other cluster leaders. In addition, discussions and joint work was initated with the real-time components cluster (partners CEA and MDH) and with affiliated partners VTEC and Volvo car, the purpose of which was to achieve a better understanding of different approaches towards model and tool integration. This topic is today addressed by many researchers and companies, spurred by the increasing product complexity and needs to support early integration of models

representing different aspects and parts of a product. Several variants of model-based approaches are today advocated to facilitate systems integration. A survey was conducted including a number of representative efforts that address multiple concerns or views including modeling languages such as AADL and EAST-ADL as well as model integration environments such as GeneralStore, ToolNet, and Fujaba.

Part of this work was carried out in connection to the new European research project, ATESST, involving KTH, Volvo (affiliated partner) and CEA (real-time components cluster partner), as well as other automotive companies.

www.atesst.org

Output from Achievement: Interactions and characterization

- An extended tool survey essentially with complementing information from Inria on the Syndex tool and from Univ. of Salzburg on the GIOTTO tool, [5].

- A jointly authored paper surveying different approaches towards model and tool integration, highlighting their commonalities and differences regarding basic integration mechanisms and engineering support, [6].

- A better understanding of the challenges, integration characteristics and types of solutions available with respect to model and tool integration.

Difficulties with Achievement: Interactions with other ARTIST2 clusters

Interactions with other clusters is resource/time demanding because it requires that disciplinary gaps (terminology and mutual understanding) are bridged. This is even more difficult today because people in academia and industry tend to be extremely busy. Therefore, dedicated efforts and resources/time have to be devoted for this purpose. The partial success reported here is due to ARTIST2 as such, already existing connections, and new projects, such as ATESST. We believe there is more potential with this type of interactions.

Achievement: Tool Integration

An example of how the to co-design tools TrueTime and Jitterbug from LUND can be combined has been developed. In [1] Truetime is used to, using simulation, derive the sampling jitter distributions and the input-output latency distributions for a controller task set executing in a real-time kernel. These distributions are then used by Jitterbug to analytically evaluate the reaulting control performance.

The tools are interfaced through the Matlab workspace. Another approach to combine the tools is for performance evaluation of nonlinear control loops. Jitterbug is able to analytically evaluate a quadratic control performance function for linear systems. If the control loop under investigation instead is nonlinear (either the control law or the controlled plant) then the same quadratic control performance can be evaluated by Truetime through simulation.

Based on the experiences of the AIDA toolset, an experimental model integration and management platform has been developed at KTH [11, 13]. Interfaces from Simulink and Dome, representing domain tools (in this case used for function and hardware desin respectively), were implemented to the platform which was also exercised with case studies on architectural design. The corresponding tool integration architecture draws upon experiences from mechanical engineering where product data management tools are used to store design information, and with interfaces and various levels of integration to design tools, e.g. for CAD and CAM. The design and implementation of fine-grained model management of functions, software and hardware turned out to be quite feasible using existing commercial PDM tools,

although a complete evaluation including performance, scalability etc. has not been carried out [11, 12].

Output from Achievement: Tool Integration

Apart from the papers mentioned in the previous paragraph, a better understanding of the problems facing tool integration has been achieved during year 1.

Difficulties with Achievement: Tool Integration

The efforts required for actual tool development and integration must not be underestimated. The progress in these areas depends to a large extent on the available research projects that provide explicit funding to these activities.

Achievement: Further development of individual tools

Further development of the tools developed by LTH, Jitterbug and Truetime, and by CTU, Torsche. The work at KTH on a model and tool integration platform was reported in the previous paragraphs.

Jitterbug: The development of a graphical user interface for Jitterbug has started. Currently, the user interface of Jitterbug is purely text-based. However, Jitterbug is based on block diagrams and state automata, two formalisms for which graphical interfaces are very natural. In the current GUI approach a graphical interactive interface has been developed in Java and Swing. In this interface the user develops the block diagram and state automaton models using mouse-based drag-and-drop techniques. When the user decides to perform a performance evaluation, the user interface models are interpreted and the corresponding text-based Jitterbug Matlab commands are created. These commands are then piped to Matlab, which runs as a compute engine executing the Jitterbug commands and returning the result. The GUI is at the time of writing currently completed to around 80%. With the GUI we expect the usability of Jitterbug to increase substantially.

TrueTime: A new version (1.4) of TrueTime has been released. The version includes support for semaphores (in addition to the already existing mutexes), and blocking mailboxes. The possibility to have user defined radio models for wireless networks has been added, as well as support for implementing ad hoc routing protocols, e.g. AODV. At the time of writing, the previous release (1.3) has been downloaded more than 1.900 times.

Torsche: The development of a simulation and implementation support for DSP applications in TORSCHE has started providing several case-studies. Further, TORSCHE has been extended by a simple response time analysis for the set of periodic tasks running under operating system with fixed priority preemptive kernel. Therefore one set of input parameters (computation times, periods, priorities) may be used to run simulation in True Time and response time analysis in TORSCHE. A new version (0.2) of TORSCHE has been released. The version includes new scheduling algorithms (Horn, List scheduling with various parameters, Scheduling with start time related deadlines, Cyclic scheduling), support for random generation of test cases, graph algorithms and interface to ILP solvers.

Output from Achievement: Further development of individual tools

- Jitterbug: http://www.control.lth.se/~lincoln/jitterbug/
- Truetime: http://www.control.lth.se/truetime/
- Torsche: http://rtime.felk.cvut.cz/scheduling-toolbox/

Difficulties with Achievement: Further development of individual tools

No difficulties encountered

Work in Year 3

Achievement: Developments of individual tools/platforms developed by cluster partners (all partners)

LTH has developed the TrueTime simulation toolbox in a number of directions. In September 2006 Version 1.4 was released. It contained the following new features:

- support for semaphores (previously only monitors were supported) and for blocking reads and writes to mailboxes
- the possibility to user-defined path-loss models for the wireless network blocks
- the addition of an AODV ad hoc routing protocol example
- improved execution speed of more than 100% for the real-time kernel blocks

In January 2007 Version 1.5 was released. It contained the following new improvements::

- major performance improvements for the network blocks
- new network interface blocks which make it possible to use the network blocks standalone, without any real-time kernel blocks, something that is of interest in certain networked control applications

Each version is downloaded by between 1,000 – 1,500 user over the course of a year (version 1.5 has been downloaded 963 times since its release in January 2007). During the year we have also learned about several new users. For example at Universite d'Evry Val d'Essonne TrueTime is used in the final project course in control. In this control the students apply all the theory they have learned in earlier courses on a distributed control example where an inverted pendulum controller is closed over a CAN network. Before the students are allowed to try their design on the real physical system, they must develop a TrueTime model that verifies that their design works. The same university have also used TrueTime in several master thesis projects with French industry, including PSA and GM. The number of Artist2 partners who are using TrueTime increases steadily. Some of these that we are aware of are Aveiro, SSSA/Pisa and TU Vienna.

A drawback with TrueTime is that it is based on Matlab/Simulink. In a master thesis project (Kusnadi 2007) we have successfully shown that it is possible to port TrueTime to Scilab/Scicos. We have also evauated that it is possible to use multi-threading to simulate the different user threads in a real-time kernel model. In the current TrueTime version multi-threading is emulated which implies that context switching also must emulated. Using the new approach would greatly simplify the implementation of the kernel blocks in future releases. It would also make it easier to port production C code into TrueTime.

In (Gonzalo et al 2007) the possibility to include TrueTime in the EJs (Easy Java Simulation) simulation toolbox has been evaluated. This has been done in cooperation with UNED, Madrid.

KTH, an ARTIST2 and Control for embedded systems partner, has been using Truetime. In work at KTH, modeling and simulation concepts for dynamically configurable systems have been studied. In particular, the use of Truetime as a basis this has been evaluated. Certain abstractions, such as memory, and modelling of dynamic configurations are not explicitly support by Truetime and are the subject of further work [Naseer et al., 2007].

Achievement: Development of TORSCHE (CTU).

The CTU tool called TORSCHE Scheduling Toolbox for Matlab has adopted several extensions and changes (see [Sucha et al (2006) and Kelbel and Hanzalek (2006)] and http://rtime.felk.cvut.cz/scheduling-toolbox/). Beside of development of new scheduling algorithms, the integration with Truetime has been shown on typical examples of DSP algorithms implemented on FPGAs where Torsche profits from cycle-exact simulation executed by Truetime. Coupling of Truetime and Torsche was demonstrated during labs of Graduate Course on Embedded Control Systems, May 7-11, 2007, Lund, Sweden. Further we finished a work on graph editor and web based production of scheduling results in Gantt charts written in Perl and Metapost.

Feature screen casts were published 15-Jun-2007. The screen casts have been developed in order to simplify the use of scheduling algorithms within Matlab environment. The screen casts focus on the following problems: How to create Scheduling objects. How to solve a optimization problem. How to work with Graph object. How to write an easy scheduling algorithm (Earliest Release Time algorithm). How to implement a simple algorithm (Minimum spanning tree).

Achievement: The TrueTime RUNES demonstrator (LTH, KTH, UIUC)

During the year most of the TrueTime development has been funded by the EU/IST project RUNES. In RUNES a large demonstrator involving wireless networks and autonomous mobile robots that are used as mobile network routers has been developed [Årzén et al 2007a,b]. In parallel with the physical scenario a large TrueTime model of the scenario has been developed [Årzen et al 2007]. The TrueTime model includes models of heterogeneous mobile robots consisting of both AVR microcontrollers and Tmote Sky “motes”, and of stationary sensor network nodes. The model further includes ultrasound-based localization and data fusion using Extended Kalman filters. The movement of the robots and the connectivity status of the sensor network are animated dynamically. The developed model is most likely one of the largest TrueTime models ever developed. Our international partner UIUC has been involved in the implementation of the TrueTime model of the AODV protocol.

Achievement: The Saint demonstrator (KTH with Scania and Enea)

The KTH Saint demonstrator (http://www.md.kth.se/saint/) and model integration and management platform. The Saint demonstrator, which has been developed in cooperation with Scania, constitutes a scale-model truck including mechanics, sensors, actuators and distributed control system. It incorporates a simple static middleware and advanced configuration environment which enables a user to configure which functions (e.g. adaptive cruise control and collision avoidance by braking) to be included. The configuration tool will then based on information stored in a product data management system, identify the corresponding software components and perform and suggest an intelligent allocation of software components to hardware nodes, build the complete (or partial) system and download it to the truck. The demonstrator illustrates the limitations of middleware approaches that do not explicitly consider real-time behaviour. It also provides a foundation for futher experiments with life-cycle model-based information management and domain tool/aspect integration encompassing mechanics, software and electronics, [Larses et al (2007), Axelsson et al (2007).

Achievement: Developments of integrated environment for embedded control systems (CTU with UNIS Ltd. and Czech Academy of Sciences)

The motivation of our work (for more details see Bartosinski et al (2006) and Bartosinski et al (2007)]) is to make a Matlab/Simulink compatible design tool for embedded control systems compliant with HIS and AUTOSAR. The tool is based on Processor Expert (http://www.processorexpert.com/), a component oriented development environment supporting several hundreds of microcontrollers, and Matlab/Simulink (http://www.mathworks.com/) which is the de-facto standard in the rapid prototyping of the control applications but it does not have an adequate HW support. The objective is to provide an integrated development environment for embedded controllers having distributed nature and real-time requirements. Therefore we discuss the advantages of using an automatically generated code in the development cycle of the control embedded software. We present a developed block set and Processor Expert Real-Time Target for Matlab Real-Time Workshop Embedded Coder. The case study shows a development cycle for a servo controller.

Achievement: Model-based embedded systems engineering (KTH, Volvo, CEA, other ATESST partners and the Univ. of Hull)

In connection to the ATESST project (www.atesst.org) KTH has been investigating model transformations between UML, Simulink and safety analysis tools. In the ATESST project, a UML profile for automotive embedded systems modelling is developed. The goal is to provide support for coherent systems level modelling, while enabling integration with domain tools.

UML-Simulink. A new UML tool environment, Papyrus, developed by CEA is used to demonstrate the concepts, but the idea is that the profile should be useful in commercial UML tools. During this year scenarios (providing motivation) and ways of performing structural and behavioural transformations between UML and Simulink have been investigated. The work has been promising but there are several outstanding research topics which will be dealt with in the subsequent work [Shi et al (2007), Cuenot et al (2007a)], one being to what extent the UML can be formalized using its available extension mechanisms and current definition of behavioural models. There are also several technological issues involved, such as how well UML tools support profiling (this is important given the interest in developing profiles).

Continuous-time modelling in SysML/UML2. Related to the above effort, KTH has been investigating the possibilities for continuous-time modelling in SysML/UML2. This type of modelling can in principle be performed on different levels of abstraction, and using different types of diagrams such as parametric and acitivity diagrams; we have performed an initial investigation of different approaches, found a limitation in the parametric diagrams and concerns when using the activity diagrams – these are subject for further work [Sjöstedt et al (2007)].

UML-Safety modelling and analysis. Safety analysis for embedded control systems, including analysis at functional as well as implementation level have been investigated. A new fault/error/failure modelling concept linked to systems level modelling is being developed, partly in cooperation with the Univ. of Hull. The approach builds on model-based embedded systems engineering integrated with extensions of classical safety analysis techniques such as failure-mode and effects analysis and fault-tree analysis [Cuenot et al (2007a), Cuenot et al (2007b)].

Achievement: Better understanding of industrial practices in automotive embedded systems model based development (KTH)

Interviews and studies of industrial practices in the area of automotive embedded systems area have been carried out. This work is motivated by the need to better understand the gap

between research and industrial practices, what the industrial challenges are, and to provided insight into how systematic approaches to model-based development can be introduced in industry. The studies have provided several important insights regarding the consideration of process and organizational constraints when introducing new methods/tools, and how model-based development supports product/process/organization integration [Adamsson (2007), Malvius (2007)].

Achievement: Higher level of interaction between clusters and industry on model-based development (all partners)

Martin Törngren of the Control for embedded systems cluster took the initiative to raise and discuss the needs of ARTIST2 actions dedicated to synchronization between various platforms, models and tools. This initiative received a positive response from the other ARTIST2 clusters, resulting in two ARTIST2 workshops during the 3rd year, one collocated with the DATE conference, and the other one collocated with the CAV conference.

http://www.artist-embedded.org/artist/-ARTIST2-Workshop-at-Date-07-.html

http://www.artist-embedded.org/artist/-Tool-platforms-for-modelling-.html

The workshops served both as dissemination and for discussing design flows, methodology, tools and modelling approaches for embedded systems (see the dissemination achievement for the links). The workshops were successful but more efforts in this direction are needed; it takes time and resources to bridge disciplinary gaps.

A larger KTH/Industry seminar was organized on Aug. 30th at KTH to discuss embedded systems challenges, industry/academia interactions, and activities of the forthcoming KTH embedded systems centre cooperation. The European commission was represented at the meeting and provided a talk on joint technology initiatives, focusing on ARTEMIS.

http://www.md.kth.se/RTC/KTH_es_seminar2007.html

Final Results

The technical achievements in terms of NoE integration include sharing of developed models (the Bridgit case studies), further development of tools to make them more easily adaptable and usable by the research community (TrueTime), sharing of tools with other partners, joint development of model integration platforms (including KTH, CEA and Volvo), joint development of automotive embedded systems middleware (including KTH, Bosch and Daimler), and information exchange through jointly organized workshops with other clusters and also Artist2-external organizations. A high level of interaction between the partners of the clusters, with other clusters as well as with industry and other universities has been maintained. These interactions are further elaborated in sections 2.3.1, 2.3.3 and 2.3.5.

Developments of TrueTime:
TrueTime is a Matlab/Simulink-based tool for co-simulation of real-time control systems that has been developed at Lund University since 1999. Using the tool, it is possible to build detailed simulation models of plant dynamics, controllers, task scheduling, wired and wireless network communication, mobile robots, etc.

TrueTime 1.6 has been in the making during 2007 and 2008 and has been released at the end of September 2008. The most important change in the new version is that the simulator software will be released under the GNU General Public License (GPL). This will allow researchers and developers to expand the functionality of TrueTime and add new features, thereby boosting the further development of the tool.

New features in TrueTime 1.6 include a new Ultrasound Network block, better support for Constant Bandwidth Server scheduling, and an improved user interface. The user no longer has to graphically connect all nodes and networks; this is resolved by the tool itself.

During the year, TrueTime has been used in laboratories and exercises in the courses and workshops given by the cluster, e.g., the Embedded Control Systems graduate course given at KTH and the workshop given in connection with the IFAC Congress in Seoul. In (Gonzalo et al 2008), it is described how TrueTime can be combined with the EJs (Easy Java Simulation) simulation toolbox and how the TrueTime way of modeling real-time kernels also can be implemented in EJs. This has been done in cooperation with UNED, Madrid. TrueTime is also the subject of a book chapter in a forthcoming CRC Press book (Cervin and Årzén, 2009).

TrueTime will also play an important role in three new embedded control system projects started in 2008. In ACTORS, TrueTime will be used by ULUND, TUKL, and SSSA to evaluate a feedback-based reservation management schemes. Here, the goal is to extend the real-time kernels in TrueTime to also support multiple cores and Linux-type weighted fair-queue scheduling. In CHAT, TrueTime will be used in the context of clock synchronization by Siemens and ULUND, and finally in WIDE, a modeling and simulation toolbox for networked control systems will be developed by KTH and others that will be based upon TrueTime.

TrueTime in EUROSYSLIB:
Lund University is also a member of the ITEA 2 project EUROSYSLIB led by Dassault Systems.

The ultimate objective of EUROSYSLIB is to make Modelica (http://www.modelica.org) the de-facto standard language for embedded system modelling and simulation. In order to support this major product lifecycle management effort, the EUROSYSLIB consortium, composed of 20 European partners, is committed to delivering a large set of high-value, innovative modelling and simulation libraries based on the freely available Modelica object-oriented modelling language. The role of Lund University is to develop a network simulation library for Modelica with features that are similar to the network blocks in TrueTime <http://www.control.lth.se/truetime>.

TrueTime in DySCAS:
KTH has been using TrueTime within the DySCAS project to model and simulate dynamically configurable systems, in particular for the evaluation of different run-time reconfiguration algorithms and quality of service management approaches, see Feng et al. (2008 - CDC) and Feng et al. (2008 - report)]. This work is also touched upon further below under the heading DySCAS middleware.

TORSCHE developments:
The CTU tool called TORSCHE Scheduling Toolbox for Matlab has adopted several extensions and changes (see http://rtime.felk.cvut.cz/scheduling-toolbox/). Scheduling is one of key factors in resource aware design of embedded control systems. Due to the complexity of various control systems executed on parallel and distributed resources, the design process requires interactive software supporting optimal scheduling of these resources and intuitive simulation of resulting schedules influencing the control performance. The methodology for code generation from a description in a subset of the Matlab language, to simulation in TrueTime may be demonstrated on signal processing applications (see Šůcha, P., Kutil, M., Hanzálek, Z., IFAC 2008). Furthermore, a case study encompassing a simulation in virtual reality toolbox has been realized in order to illustrate the architecture and capabilities of the tool-chain. In particular, the release from 12-Oct-2007 includes new scheduling algorithms (McNaughtons's algorithm, Hu's algorithm, new version of algorithm for cyclic scheduling,Coffman's and Graham's Algorithm,

Bratley's algorithm and scheduling algorithm for problem P|rj,prec,~dj|Cmax, P||Cmax, 1||sumUj, 1|rj|Cmax) graph theory algorithms (Dijkstra's algorithm, Algorithm for spaning tree problem, Tarjan's algorithm, Graph coloring). TORSCHE is distributed with the book by Michael Pinedo: Scheduling: Theory, Algorithms, and Systems (Third Edition).

The DySCAS middleware architecture and development framework (KTH with Volvo, Enea, Bosch, Daimler and Univ. of Greenwich)

DySCAS is an autonomic platform-independent middleware for automotive embedded systems which is being developed in a European project (www.dyscas.org). The concepts and architecture are motivated by the need for a higher flexibility and automatic run-time reconfiguration in embedded systems (moving from static to dynamic configurations). Examples of situations that may require reconfigurations of how software is allocated to hardware and how resources are managed include the addition of new functions (e.g., through software upgrade or by attaching new hardware devices) to an existing embedded systems, and a scenario where the embedded system itself detects poor performance or failing components.

In DySCAs, self-management is achieved in terms of automatic configuration for context-aware behavior, resource-use efficiency, and self-healing to handle run-time detected faults. The self-management is governed by the use of policies distributed throughout the middleware components.

The status of the project is that a system architecture has been developed and that reference implementations are being developed [Chen et al. (2008 - ERTS), Anthony et al. (2008)].

In order to develop a system such as DySCAS, various types of tools are required covering a wide range of design activities from architecture modeling and design, verification/validation through simulation and formal analysis, to prototype implementation. Clearly, for DySCAS-type systems to make it to the market there will be a need for tools supporting the entire system life cycle including tasks such as configuration, deployment and maintenance, however such tools are not treated here. The following figure illustrates current tools (solid line boxes and connections) and future possible tools (dashed line boxes and connections) and their interconnections. Current support for the design includes simulations, safety analysis using FMEA and formal verification with the purpose of providing feedback to the architecture modeling and specification work which is carried out using UML. In general, modeling and analysis of these different aspects plays the important role of improving our understanding of DySCAS systems, in communicating DySCAS concepts by means of models, and in verifying and validating designs. In the following, current support tools and experiences are described (solid line boxes).

The architecture has been captured by UML models encompassing structural and behavior aspects. Model transformations have been defined but so far not automated in order to transfer architectural models from UML to Simulink (for simulation purposes) and to XPCT[2] for formal verification (Note: The computation tool adopted in this investigation is XPTCT. While this software tool is intended for supervisory control design of discrete-event systems modeled as finite automata, we use it only for automaton computation.)

Simulation within the Matlab/Simulink environment has been carried out at two levels of abstraction:

[2] Free to download from http://www.control.utoronto.ca/DES.

- Logical simulations. The main purpose of these models and simulations is to verify structural/architectural properties (component model, interfaces and signals) and logical behavioral properties (state-machines and activity diagrams). The modeling emphasizes platform independent modeling and supports both interface verification as well as validation through simulation of the system behavior. We have chosen Matlab/Simulink/SimEvents as tools due to their ability to support simulation of discrete event and continuous time systems simultaneously.
- Base simulations. The main purpose of these models and simulation is used to evaluate system behaviors, including algorithms for configuration management, quality of service and load balancing. The modeling has included explicit platforms abstractions in order to incorporate aspects such as allocation to processors, platform performance as well as application performance. For these simulations we have chosen the TrueTime (Ohlin et al., 2008) toolbox due to its support for modeling logical as well as real-time operating systems and network protocols.

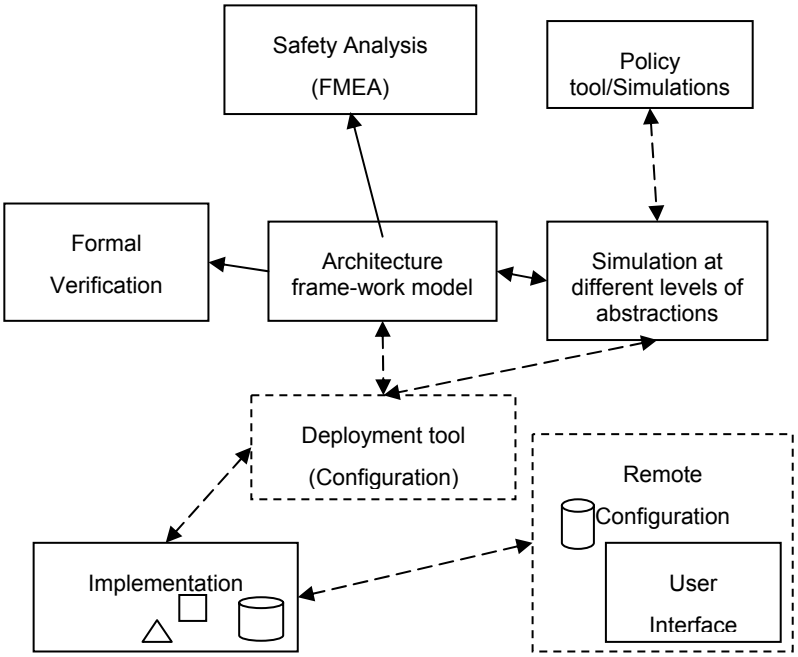

Existing and desired design activities (methods/tools) used in the development of the DySCAS architecture. The solid lines illustrate where information has been transferred by manual model transformation. In the future we wish to automate these transformations. The dashed lines illustrate desired new tools and connections, for example where a configuration tool can be used as a basis for defining a concrete DySCAS implementation and for configuring simulations.

Continued work on the KTH Saint demonstrator (http://www.md.kth.se/saint/).

As part of the DySCAS project, the Saint platform is reused since it well represents statically configured systems as they appear in the automotive industry today. While the Saint middleware is simplified, it is conceptually similar to the AUTOSAR standard. DySCAS systems will probably first be introduced for less critical and fastly changing functions, represented by the Infotainment and Telematics domains. This means that a DySCAS system must be possible to interface to the other domains inside a vehicle. In addition, a migration strategy from statically to dynamically configurable systems has to be considered.

The main approach taken within DySCAS has been to investigate the use of a gateway approach to separate the two networks (Garcia, 2008). A gateway approach is a natural solution since gateways are frequently used today to separate the existing domains. This approach has several advantages and allows the two networks to coexist. Integrating the two types of middlewares into the same domain would be a very difficult challenge

Model-based embedded systems engineering (KTH, Volvo, CEA, other ATESST partners and the Univ. of Hull)

In the ATESST project, a UML profile for automotive embedded systems modelling is developed. The goal is to provide support for coherent systems level modelling, while enabling integration with domain tools. During the 4th year of ARTIST2, the ATESST project finished and the follow up project, ATESST2, started (www.atesst.org). The work reported last year has continued.

UML-Simulink. A new open-source UML tool environment, Papyrus (www.papyrusuml.org), developed by CEA is used to demonstrate the concepts. A mapping between Simulink behavioural models (focusing on discrete-time and continuous-time systems) to UML composite structure diagrams AND activity diagrams has been developed. The mapping, which is valid in both directions, has the interesting property to provide an explicit representation of Simulink structure and behaviour. The behaviour of a Simulink model is captured on one hand by the algorithms, the types of blocks and their connections, and on the other hand by the Simulink simulation engine and its ordering of the blocks to provide an operational model for simulation. An algorithm has been developed that takes these aspects into account and produces UML models from the Simulink model. A partial tool implementation is currently available and work towards a full realization is ongoing, [Sjostedt et al (2008)].

UML-Safety modelling and analysis. Safety analysis for embedded control systems, including analysis at functional as well as implementation level have been further investigated. A new approach for representing safety cases based on the previously developed error modelling has been developed. The key solution in the approach is to provide an integrated information model in terms of the EAST-ADL architecture description language. In system modelling it then becomes possible to link related information entities (requirements, software-hardware architectures, error models, hazards etc.) and related analysis tools (safety, simulation, etc.), [Chen et al. (2008 - Safecomp), Törner et al (2008)].

Web-based tools for embedded systems analysis

At UPV, some web-based tools to design and analyse the dynamic memory use by embedded applications, and for the analysis of the scheduling of hierarchical real-time systems have been developed. The hierarchical aspects concerns the use of a so called “hypervisor” approach that allows several real-time operating systems to run concurrently on the same processor.

- http://wks.gii.upv.es/schedtools/
- http://wks.gii.upv.es/stools

Adaptive Real-time, HRT and Control (NoE Integration)
Led by Karl-Erik Årzén (Lund University)

Partner teams (leaders): Karl-Henrik Johansson – KTH (Sweden), Anders Robertsson – LUND (Sweden), Karl-Erik Årzén – LUND (Sweden), Alfons Crespo – UPVLC (Spain), Pedro Albertos – UPVLC (Spain), Martin Törngren – KTH (Sweden), Giorgio Buttazzo – (SSSA, Pisa), Albert Benveniste – (INRIA), Gerhard Fohler – (Univ Kaiserslautern), Werner Damm – (OFFIS), Vladimir Kucera – (CTU), Zdenek Hanzalek – (CTU), Hermann Kopetz (TU Vienna), Luis Almeida - (University of Aveiro), Eduardo Tovar - (Polytechnic Institute of Porto).

Affiliated teams (leaders): Johan Eker – Ericsson (Sweden), Lui Sha - University of Illinois (USA), Tarek Abdelzaher - University of Illinois (USA), Pau Marti – (UPC), Juan Antonio de la Puente - (UP Madrid).

Overview: The situation at the beginning of this activity, i.e., at the start of ARTIST2 was the following. A number of projects already allowed some structuring and interaction, e.g., FLEXCON (Swedish national project) Flexible Embedded Control Systems involving Lund and Mälardalen), OCERA (European project) Open Components for Embedded Real-Time Applications, involving Pavia, UPVLC and CTU, RECSYS (European project) Real-Time Embedded Control of Mobile Systems with Distributed Sensing involving KTH, ARTIST (European Accompanying Measure) Advanced Real-Time Systems involving a majority of the partners, and FIRST (European project) Flexible Integrated Real-Time Systems Technology involving Pavia (Pisa) and Mälardalen.

There also existed strong links between the core partners and the affiliated partners, e.g., between Lund and Ericsson, between Lund, Virginia and Illinois, and between UPC and Mälardalen.

Work in Year 1

Work achieved in the first 6 months of Y1

- Kick-off meeting held
- UPVLC (Crespo) has evaluated the performance of the scheduling policies related to offer constant bandwidth behaviour. In conjunction with SSSA (Lipari), a new version of the CBS called IRIS was developed. This new algorithm was implemented and evaluated in a real-time environment providing both hard and soft real-time constraints. The IRIS algorithm was implemented in RTLinux and included in the distribution of the OCERA project.
- Collaboration between Mälardalen (Fohler) and LUND (Cervin) about the combination of the jitter margin index and flexible scheduling methods.
- CTU studied holistic scheduling methods and analyzed a case study using the MAST tool (Cantabria).
- Karl-Erik Årzén contributed to the Artist2 workshop on diagnosis in Vienna, Dec 20-21, organized by the HRT cluster

Work achieved in months 6-12 of Y1

- The Lund Workshop was held. The interaction between the control cluster, the participants from the ART cluster and the US affiliated partners was very valuable
- In order to add flexibility to the real-time applications UPVLC has developed a nanokernel called Xtratum. Xtratum is a thin layer of software that provides a simple

and convenient API to access interrupt mechanism and timer devices. Xtratum permits the execution of environments/applications spatial and temporal isolated. Xtratum has been developed under the OCERA project.

- Collaboration between LUND (Cervin) and Pavia (Buttazzo) about the use of the Shark real-time kernel as a shared platform for implementing control applications.
- Collaboration between LUND (Cervin) and Ericsson (Eker) on distributed versions of the control server model
- A collaboration between UPC (Marti), Mälardalen (Fohler) and LUND on feedback scheduling of control system has been initiated
- CTU has built up several demonstrators for communication components based on the OCERA architecture (UPVLC, SSSA, CTU) including fish breeding control and supervision system (process control application), remote programming of mobile robot (robotics and supervision), human machine interface for autogiro (data acquisition and visualization), robotic arm demonstrator (servo control).
- Interaction between CTU and Aveiro (Almeida) on deadline constrained scheduling on FPGAs and multicast traffic optimization.

Work in Year 2

Achievement: Organization of Workshop

The workshop Interaction between control and embedded electronics in automotive industry was jointly organized by the RT Components and the Control clusters in Innsbruck, March 23. It was co-located with the Beyond AUTOSAR meeting organized by the network activity "Forums with Specific Industrial Sectors". Three invited presentations were given by Stefan Kowalevski (RWTH Aachen), Karl-Erik Årzén (Lund University), and Carlos Canudas de Wit (LAG Grenoble) followed by a panel discussion. A more detailed description of the content and focus of the presentations is given in the activity report of the "Forums with Specific Industrial Sectors" activity. Several conclusions can be drawn:

- There is a permanent misunderstanding between control & software engineers in the automotive industry
- Regarding the relative merits of ET/TT, control design aspects provide complementary views, not considered in the embedded design community. For example, in general a long but constant controller input-output latency is worse from a control performance point of view than a shorter but time-varying latency, also if the former constant latency is taken into account in the control design.
- The control systems in automotive systems are often structured in a multi-layer or multi-cascade fashion. This further increases the need to minimize the input output latency and puts special requirements on component-based architectures. For example, it is important to organize the computations in such a way that first only the parts of the controller components that are needed for the generation of the component outputs are calculated and then, afterwards, the parts of the components that are responsible for updating the state of the controller components are calculated. This is something that is well-known within the field of process automation, but for some reason has not yet spread to, e.g., the automotive systems area

- In an automotive system there is only a limited amount of sensors and actuators. Both the sensors and actuators are typically used by several control systems or control functions. In an integrated system it is important to make it possible for several functions to use the same physical sensors and actuators, rather than, e.g., use several sensors to measure the same physical entity, something which is not uncommon in federated architectures. Hence sensor and actuator component should have a special role in a component-based automotive system.
- Today, the structure of the control systems in an automotive system is to a large degree derived from the constraints of the federated system architecture. In an integrated system new possibilities for structuring the control systems open up. Hence, it would be worthwhile to take a completely new look upon how the overall control system for a car ought to be structured, including powertrain control, chassis control, safety systems, etc.

Achievement: Joint Research Activities Involving the ART and the Control Cluster

The joint research initiatives that were started during Y1 have continued. These include

- Anton Cervin (Lund) and Giorgio Buttazzo (Pisa) have worked on a comparison of jitter reduction techniques for control tasks. When implementing a controller in a multitasking operating system, there is a risk that the control loop will experience delay and jitter due to pre-emption from other tasks. Several jitter control methods have been proposed in the literature, and they all have different strengths and weaknesses with respect to timing and control performance. In this work, they have compared and evaluated four different task models: the Standard Task Model (STM), Reducing Jitter by Task Splitting (RJTS), Reducing Jitter by Advancing Deadlines (RJAD), and Reducing Jitter by Non Preemptive Execution (RJNP). It is found that RJTS is good for jitter reduction, but introduces a long delay which gives sluggish control performance. RJAD works well for reducing both jitter and delay, and gives good control performance in most cases. RJNP reduces input-output jitter to a minimum but may cause some tasks to miss their deadlines. A conference publication describing this joint work is under preparation and a technical report is available.
- Lund (Cervin) and Pisa (Bini) have worked on optimal period selection for multiple controllers under fixed-priority scheduling. Traditionally, when scheduling controllers, it has been assumed that the deadline of each control task is less than or equal to its period. Under fixed-priority (FP) scheduling, this typically implies that the processor cannot be fully utilized. In this work, they have explored what control performance is possible to gain by moving outside the FP schedulability bound. Utilizing a simple upper bound on the response time of a task, the input-output delay can be bounded. Combining this bound with an approximate expression for the control performance (as a function of the rate and the delay of the controller), the optimal task periods can be found by solving a constrained optimization problem. For certain simple cases, exact analytical solutions can be found. A publication describing this joint work is under preparation.
- UPC (Marti, Selga) and Lund (Henriksson, Cervin) have worked on feedback-based scheduling of linear controllers with varying disturbance intensities. In previous work from Lund on feedback scheduling of linear controller tasks, it has been assumed that the amount of disturbances entering the control loops is constant over time. In [1] the initial states of the controlled plants are taken into account by the feedback scheduler by including the initial state in the cost function. The motivation for this is that a plant with a large error should receive more resources in order to better

cope with the disturbance. However, in all but extreme cases it is the expected future disturbances that completely dominate the cost function. In this work, they have explored how one can obtain a more reactive feedback scheduler by estimating the amount of noise in the various control loops. They have also extended the cost functions to take a constant delay (obtained using Control Servers) into account. The project has included a PhD student visit from UPC to Lund: Rosa Castañe spent 5 months (from August 2005 to December 2005) in Lund. In addition, several working meetings have taken place during 2006, in Pisa, March 2006 and Dresden, June 2006.

- Lund (Cervin) and Mälardalen/Univ Kaiserslautern (Moris, Isovic, Fohler) have continued the work on flexible scheduling of controllers based on the jitter margin. The work combines two previously developed tools and techniques for flexible real-time systems: the jitter margin and the slot-shifting algorithm. Using the jitter margin, it is possible to guarantee a level of a performance of a controller, given bound on the worst-case input-output jitter. On the other hand, the slot-shifting technique can be used to allow sporadic tasks to execute at the cost of more jitter for the periodic tasks. In this work, an off-line design method based on simulated annealing has been developed that tries to find an optimal schedule such that all control tasks meet their performance specifications, while at the same time allowing as many sporadic tasks as possible to execute. The work has resulted in a Master Thesis which recently received the price for the best Swedish Master Thesis in the field of Real-tIme and Embedded systems during 2005-2006.
- Several of the groups have focused their activities on the SHARK kernel and the TrueTime tools as common platforms for feedback-based scheduling work. In Lund a project has started in which the suitability of using SHARK in control laboratories will be investigated. UPC has modified the Truetime simulator to better study new feedback scheduling theoretical results. UPC has also added new features to Shark to allow easy implementation of feedback scheduling.
- A strong research connection is currently being established between CTU and UPCLC in the Control cluster and UCantabria, Pisa, and UYork in the ART cluster. This is funded through the FRESCOR project. Here several activities are currently being initiated, e.g., the implementation of contract-based kernels for embedded systems. Both CTU and UPVLC also participated in the ARTIST2 requirements workshop (Paris June 16 2006).

Achievement: Joint Summer School

The summer school First European Laboratory on Real-Time and Control for Embedded Systems was organized in Pisa, Italy, July 10-14, 2006. The number of participants was 40. http://www.artist-embedded.org/FP6/ARTIST2Events/Events/RT-Control/

Work in Year 3

Achievement: Organization of 2nd International Artist2 Workshop on Control for Embedded Systems

The 2nd International Artist2 Workshop on Control for Embedded Systems was successfully organized in Urbana-Champaigne at Univ of Illinois with Tarek Abdelzaher as the local host. The topics of the workshop were Real-Time and Control in Sensor/Actuator Networks, Control in Cyber-Physical Systems, Event-Based Control and Computing, and Control of Software Errors. The workshop activity was intended and planned as a network activity. However, due to

two unfortunate late cancellations from participants from the ART cluster, we ended up with Artist2 participants representing only the control cluster (Lund, KTH, UPVLC, and the international associated partner UIUC). However in spite of this the workshop was very valuable.

The number of participants was 20, excluding PhD students. Out of these six were form Europe. The US control community was represented by Bruce Krogh, Geir Dullerud and Michael Lemmon. The US real-time systems community was represented by Lui Sha, Tarek Abdelzaher, Marco Caccamo, P.R. Kumar, and Chenyang Lu. Industry was represented through Microsoft and PARC. The conclusions from the workshop are available as a separate document available through: http://www.artist-embedded.org/artist/-Control-for-Embedded-Systems,810-.html

Achievement: Joint Research Activities Involving the ART and the Control Cluster

The joint research initiatives that were started during Y1 have continued also into Y3. These include:

- The joint work between Lund and SSSA/Pisa on jitter reduction methods for control system applications has continued and has been extended to also include Halmstad University.
- TUKL and Lund have recently started working on a control-based evaluation infrastructure for combined offline and online scheduling. The Truetime tool of Lund is used as an initial basis for this work.
- Aveiro and UPC are working together on dynamic rate and control adaptation in networked control system. The TrueTime tool from Lund is used as the simulation platform for this work. In particular they study dynamic rate adaptation in simple microcontroller-based computer control systems, maximizing the admitted load in such systems while minimizing the control performance degradation. They also work on minimizing the resources (computing and communicating) consumed by a feedback control loop using a dual-rate (dual-controller) approach, i.e., one nominal controller and a low bandwidth one, for periods of near stationarity.
- CTU and Porto have worked together on simulation of wireless radio protocols. The IEEE 802.15.4 protocol has the ability to support time-sensitive Wireless Sensor Network (WSN) applications due to the Guaranteed Time Slot (GTS) Medium Access Control mechanism. Recently, several analytical and simulation models of the IEEE 802.15.4 protocol have been proposed. Nevertheless, currently available simulation models for this protocol are both inaccurate and incomplete, and in particular they do not support the GTS mechanism. An accurate OPNET simulation model, with focus on the implementation of the GTS mechanism has been proposed. The motivation that has driven this work is the validation of the Network Calculus based analytical model of the GTS mechanism that has been previously proposed and to compare the performance evaluation of the protocol as given by the two alternative approaches. Additionally, and probably more important, based on the simulation model they proposed a novel methodology to tune the protocol parameters such that a better performance of the protocol can be guaranteed, both concerning maximizing the throughput of the allocated GTS as well as concerning minimizing frame delay.
- UPC and Lund are working jointly on feedback scheduling of real-time control tasks. A theoretical framework for feedback scheduling of real-time control tasks was reported in last year's deliverable. This year the work has focused on implementing a case study at UPC to show the validity of the theoretical approach. The experimental implementation corroborates the simulated results reported last year.

- A strong research connection is currently being established between Ericsson, Lund, SSSA/Pisa, TU Kaiserslautern, and Evidence though the new FP7 STREP project ACTORS (Adaptivity and Control of Resources in Embedded Systems) that recently has been approved. The project also contains two non-ARTIST2 partners and has Xilinx US as an associated partner. The goal of the project is to combine reservation-based scheduling, data-flow actors-based programming models, and feedback control for control and media processing applications on Linux-based multicore platforms and FPGA platforms. The project will formally start 1 Dec 2007, but the preparations for the project started already in February 2007.
- During the HSCC conference in Pisa in April 2007 a special meeting was held involving several partners in this activity, on the topic Future trends on Networked and Embedded Control Systems. The participants were Michael Lemmon (University of Notre Dame, Indiana, USA), Paulo Tabuada (University of California at Los Angeles, California, USA), Giorgio Buttazzo (University of Pisa, Italy), Enrico Bini (University of Pisa), Anton Cervin (University of Lund), Manel Velasco (Technical University of Catalonia) and Pau Martí (Technical University of Catalonia). After reviewing the state-of-the-art on Networked and Embedded Control Systems, several topics were discussed. One of the clear conclusions was that event-based scheduling, or alternatively event-based control, is a major trend that should be further studied because it has the ability of dramatically saving computing resources while guaranteeing acceptable control performance.

Final Results

Achievement: Optimal period selection for multiple controllers under fixed-priority scheduling (LUND, SSSA/Pisa)

When several digital controllers should execute on the same platform, a key question is how to distribute the computing resources among the control loops. For controllers based on periodic sampling, this boils down to selecting sampling periods for the set of controllers. In this work, they are interested in a static design problem, where the objective is to assign sampling periods in order to minimize the aggregate cost (performance index) of the controllers [5]. For controllers, the delay between sampling and actuation has a great impact on the performance. For this reason, it is important to include a model of the delay in the design problem. They have developed an approximate response-time analysis that allows them to estimate the delay for a controller under fixed-priority scheduling. Assuming cost functions that are linear in period and delay, the approximate analysis allows giving analytical expressions for the optimal sampling periods. A case study with random second- and third-order plants are used to evaluate the improvements compared to methods that do not take the delay into account.

Achievement: Optimal period selection and scheduling for multiple distributed controllers (LUND, Linköping Univ)

Similar to the previous topic, this work also considers the implementation of multiple controllers on a shared implementation platform. The model used here is however more general: the system consists of a number of nodes, connected by a communication bus. The nodes and the bus may be scheduled statically or dynamically (using priorities). Again, the objective is to minimize the combined cost of the controllers in the application. The design problem is solved using a genetic algorithm that decides the execution pattern of the control tasks. In the dynamic scheduling case, simulation is used to estimate the average delay and jitter of the control tasks. The Jitterbug toolbox from Lund] is used to evaluate the control performance,

taking the delay and jitter into account. A case study with several benchmark plants is used to evaluate the performance compared to previous, heuristic design methods [6].

Achievement: Dynamic memory management (UPCLC, York)

Here UPVLC and University of York collaborate on dynamic memory management. Dynamic memory management can be considered in the proper way in real time applications due to the availability of dynamic memory allocators with constant allocation and deallocation. This work proposes a framework for handling dynamic memory in real-time systems [16]. The framework provides both a flexible contract negotiation model to adapt the memory allocated to a set of task as closed as the required during the execution.

Achievement: Event-Driven Control, Embedded Control, and Feedback Scheduling (UPC, SSSA/Pisa, Aveiro)

The group led by Pau Martí at UPC has been doing research on event-driven control, embedded control systems and feedback scheduling with strong, but informal connections to LUND. In event-based control they focuses on studying whether it is possible to derive sampling intervals for certain types of event-driven controllers. It is interesting to solve the problem because this will permit to obtain their resource demands, and ultimately, it will facilitate their schedulability analysis [2]. Here UPC collaborates with SSSA/Pisa. In embedded control the focus is on implementation of research results on low-cost microcontrollers with small real-time kernels. Here UPC collaborates with Aveiro [7]. UPC also studies practical solutions to the jitter problem. In particular, a task model including an asynchronous observer and a synchronized output operation is shown to be very effective when sampling intervals are not constant [3]. Finally, UPC studies feedback scheduling (or optimal sampling period selection) for real-time control tasks: this research has evaluated in simulation diverse existing results on feedback scheduling. The study concludes that on-line period assignment as a function of the controlled plant dynamics is the best approach [4]. However, it is also noted that jitter may be a problem for these approaches.

Achievement: A Simulation Model for the IEEE 802.15.4 protocol (CTU, Porto)

The joint work on simulation models for the IEEE 802.15.4 protocol by CTU and Porto has resulted in the two joint publications [8, 9].

Achievement: Loosely Time-Triggered Architectures (PARADES, INRIA, Trento, VERIMAG)

In joint work between INRIA, PARADES, and VERIMAG, the LTTA architecture (Loosely Time-Triggered has been proposed in the form of a middleware with associated mathematical MoCC and services. This distributed architecture, used in particular by some aircraft manufacturers, is compliant with the ARINC-653/AFDX architecture for Integrated Modular Avionics.

Kopetz' TTA has shown the way to develop embedded systems based on a comprehensive architecture formal model; being very strict, TTA has a cost in terms of resource and development time. Loosely Time-Triggered Architectures (LTTA) aim at coping with these difficulties, by relaxing the strict synchrony constraint of TTA. In a LTTA, the system is composed of Computing Units (ECU) and Communication Units where each unit has its own clock, clocks are not synchronized, and communication is by Sampling (CbS) meaning that it behaves like a shared memory with asynchronous writes and reads, see figure below.

During the reporting period we have developed two variants of LTTA:

a) By borrowing ideas from *elastic circuits* by Cortadella et al. and *latency insensitive design with back-pressure* by Carloni et al., both from circuit design area. Here the

aim is to preserve the Kahn Process Network semantics of the specification. In this approach, logical time dominates over physical time and no assumption is needed regarding the local clocks of the subsystems.

b) By developing a physical time based approach, where bounded relative drift is assumed between local clocks, an original protocol guarantees preservation of specification semantics. This protocol does not use back-pressure nor any token based mechanism.

a) Activities on approach (a), by all participants

PARADES, INRIA, and Verimag, jointly with UC-Berkeley and Cadence Berkeley, have developed approach (a), resulting in publications [10] and [11]. The approach assumes a single-clocked synchronous specification – single-clocked is not really a restriction in this context as it can be relaxed by using the extra symbol *nil* meaning the absence of a certain data at a given reaction. It is known that such specifications can be seen as a Kahn Process Network with bounded buffers. This observation has been the basis for the development of so-called *elastic circuits* by Cortadella et al. and *latency insensitive designs with back-pressure* by Carloni and Sangiovanni-Vincentelli in the area of circuit design. These are circuits with token based mechanisms. Controlling buffer overflow is achieved by implementing backward tokens controlling the permission to write in buffers – hence the term of *back-pressure.* This idea has been adapted to our case where neither writing nor reading can be blocking, see the figure above. The idea is to replace blocking by skipping. Performance of such architectures is classically studied by means of Marked Graphs, a simple form of Petri nets where Max-Plus algebra applies. Pros and Cons of this approach are:

- Pros: no assumption on local clocks; very adaptive, scales up easily to complex systems; easy upgrade.
- Cons: need for back-pressure, which results in additional links, resulting in additional requests for fault tolerance mechanisms.

b) Activities on approach (b), by Verimag and INRIA

Verimag and INRIA have further developed initial work by Verimag on the study of Airbus system architecture for low level flight control. They have come up with the systematic idea of replacing token based mechanisms by the use of purely local counters with no additional link. Each unit maintains a local counter based on its own independent clock. This local counter controls the right to acquire new input data from the communication media, perform computation steps, and write output data to the communication media. This approach is entirely local. Pros and Cons of this approach are:

- Pros: no back-pressure, no additional communication link, no blocking communication; this simplifies the design of fault-tolerance and degraded modes.
- Cons: uses boundedness assumptions on the relative drift between local clocks (the management of the local counters depends on these bounds). This means a higher cost compared to approach (a) when re-design is needed.

Control over loosely time-triggered networks has been studied in [1].

Achievement: Feedback Reservation Mechanisms in Multimedia Streaming (ART and Control cluster partners)

A number of partners from the Control and the ART cluster (Ericsson, LUND, TUKL, SSSA/Pisa, Evidence) collaborate within the FP7 STREP project ACTORS on Adaptivity and Control of Resources in Embedded Systems (http://www.actors-project.eu) that started in February 2008. In the project three key technologies are combined; reservation-based resource management, feedback scheduling and dataflow modeling of embedded media applications. The target platform is an ARM 11 multi-core platform. In the project, which is coordinated by Ericsson, three demonstrators will be developed: a cellular phone demonstrator, a control system demonstrator, and a high-performance video demonstrator. The collaboration among the partners initiated through Artist was very important for the successful outcome of the ACTORS proposal.

A coordination meeting between ACTORS and Artist2 was held in Lausanne in January 2008. Also, ACTORS and Artist2 co-organized a one day workshop on dataflow modeling for embedded systems in Pisa, May 5, 2008 (http://www.artist-embedded.org/artist/DataFlow-Modeling-for-Embedded.html).

Collaboration between the ART and the Control cluster also takes place within the ongoing previously reported FP6 project FRESCOR involving CTU, UPVLC in the Control cluster and UCantabria, Pisa, and UYork in the ART cluster. There are also strong links between FRESCOR and ACTORS.

Achievement: Wireless protocols for automation and control (KTH and PARADES)

KTH and PARADES have done joint work on minimum coding in CDMA networks and on Breath, a self-adapting protocol for wireless sensor networks in control and automation [12,13,14,15]

Control in real-time computing (Cluster Integration) Control
Led by Karl-Erik Årzén (Lund University)

Partner teams (leaders): Professor Karl-Henrik Johansson – KTH (Sweden), Anders Robertsson – LUND (Sweden), Karl-Erik Årzén – LUND (Sweden), Alfons Crespo – UPVLC (Spain), Martin Törngren – KTH (Sweden)

Affiliated teams (leaders): Johan Eker – Ericsson (Sweden), Lui Sha - University of Illinois (USA), Tarek Abdelzaher - University of Illinois (USA)

Overview: Before this activity started the different groups performed individual research on applying control-based approaches to embedded and real-time systems, e.g., feedback scheduling of servers, feedback scheduling of control systems, and control-based approaches in networking. The research area was strongly dominated by US research groups.

Work in Year 1

Since this is a rather new research area it was decided at the beginning of year1 that the main activity should be the creation of a research roadmap. The aim of the roadmap was to chart the area, provide a common platform for the coming work, and to identify the most important research directions.

- The first version of the roadmap was completed.

- A new feedback scheduling method was developed for control loops by Dan Henriksson and Anton Cervin (LUND). A paper was presented at the CDC-ECC'05 in Sevilla - LUND
- An international workshop in Control for Embedded Systems was held in Lund with 20 participants. The international affiliates Lui Sha and Tarek Abdelzaher participated and gave valuable input. A separate research agenda for the work within Artist2 was written as the output from the workshop.
- Karl-Erik Årzén and Anders Robertsson were invited to participate as the only non-US participants at a workshop on the future of control of computing systems organized by NFS and held at IBM, May 3-4, 2005
- KTH has been working on control-based error-correction in packet-switched networks, on the use of radio network feedback to improve TCP performance over cellular networks, and on network state estimation.

Work in Year 2

Achievement: Dissemination of Roadmap Material

The conclusions from the roadmap developed during year1 were summarized into a conference paper that was presented as an invited talk at FeBID'06, the First International Workshop on Feedback Control Implementation and Design in Computing Systems and Networks that was organized in Vancouver in April 2006. An extended version of this paper has also been published in the ACM SIGBED Review. The creation of the FeBID workshop series can potentially be very important for the future development of the area. The follow-up workshop FeBID'07 will be organized in Munich in May 2007, with a member of the Lund group as a technical co-chair and with another member of the Lund group in the IPC. (constituting two of the only three European members of the organizing committee – compared to 26 members from the US!).

Achievement: Control of Server systems

Control of server systems is the subject of research in Lund and University of Illinois. Lund is working on improved models for feed-forward based queuing control systems and on providing reservation-based scheduling in Linux systems using the nice value as the control signal. A natural application for the latter is web servers. The work at University of Illinois is focused on content distribution, adaptive rate allocation, and delay control. Dan Henriksson from Lund is spending the year 2006-2007 as a postdoc at University of Illinois working with Tarek Abdelzaher. The new model types derived for queuing control are also applied to traffic flow control in collaboration between CTU and LUND.

In a complementary activity at KTH, the automatic control group has been investigating distributed resource allocation mechanisms for large-scale server clusters. Optimal off-line solutions and high-performing distributed heuristics have been developed and evaluated in detailed system-level simulators of the Chameleon architecture.

Achievement: Feedback Scheduling of Control Systems

In our previous work on feedback scheduling of linear controller tasks it has been assumed that the amount of disturbances entering the control loops is constant over time. The initial states of the controlled plants are taken into account by the feedback scheduler by including the initial state in the cost function. The motivation for this is that a plant with a large error should receive more resources in order to better cope with the disturbance. However, in all but extreme cases

it is the expected future disturbances that completely dominate the cost function. In [6], we have explored how one can obtain a more reactive feedback scheduler by estimating the amount of noise in the various control loops. We have also extended the cost functions to take a constant delay (obtained using Control Servers) into account. This work has been performed in collaboration with UPC.

Achievement: Control of Communication Networks

The automatic control group at KTH has been working on theory and engineering principles for cross-layer optimization of wireless networks. Specific achievements include a theoretical framework for self-regulating protocol design, as well as detailed resource control strategies for specific network technologies. The KTH group has also worked on on-line error control adaptation in networked applications, feedback-based error-correction in feedback-based networks, stability of window-based queue control with applications to mobile terminal download, models for network congestion control, and distributed consensus algorithms.

Work in Year 3

Achievement: Control of Server systems (LUND, UIUC)

Control of server systems is the subject of research in Lund and University of Illinois. Lund has continued its work on improved models for feed-forward based queuing control systems and on providing reservation-based scheduling in Linux systems using the nice value as the control signal. A natural application for the latter is web servers. A feedback–based prediction scheme for controlling the response times in a single server queue has been investigated. This control structure has the benefit over other previously suggested control structures that no measurement of the required work of each job is needed. However, the new solution maintains the same attractive properties, regarding average response time and variance as previously suggested solutions.

The work at University of Illinois is focused on content distribution, adaptive rate allocation, and delay control. Dan Henriksson from Lund has spent the year 2006-2007 as a postdoc at University of Illinois working with Tarek Abdelzaher.

In parallel with this KTH has continued their work on control of server farms.

Achievement: Feedback-Based Resource Management in Cellular Devices (Ericsson, LUND)

Lund and Ericsson have received funding from the Swedish funding agency VINNOVA for the joint project “Feedback Based Resource Management and Code Generation for Soft Real-Time Systems”. The project will provide funding for one researcher from Lund University and one researcher from Ericsson over three years. A related EU FP7 STREP project coordinated by Ericsson has also been approved. This will, however, be reported in the activity report for the NoE integration activity Adaptive Real-Time, HRT and Control.

Achievement: Control of Communication Networks (KTH, Ericsson, ABB)

The networked control group at KTH has continued their research on analysis and synthesis of networked control systems, including resource allocation, traffic control and routing for wireless networks and distributed control and estimation. The work includes laboratory implementations and testing as well as industrial dissemination trough collaborative projects with ABB, Ericsson, Scania etc. KTH participated in the 2nd International ARTIST Workshop on Control for Embedded Systems in Urbana-Champaign, where samples of these results were presented.

Achievement: Adaptive Resource Management in Wireless Networked Embedded Systems (KTH, LUND, Ericsson)

Resource control in wireless networked embedded systems has been the subject for the collaboration between KTH, Lund and Ericsson funded by the RUNES IP during Year 3. A challenging demonstrator scenario has been developed and implemented involving mobile robots and sensor network nodes, in which control and localization techniques closed over a wireless network were combined with feedback-based radio transmit power control

Achievement: Dynamically Configurable Automotive Embedded Systems (KTH, Volvo)

Within the project Dyscas (www.dyscas.org), KTH in cooperation with European automotive industries (including ARTIST2 affiliated partners Volvo, and DaimlerChrysler), Enea (providing the OSE real-time operating systems), Univ. of Paderborn and Univ. of Greenwich, have spent considerable efforts on dynamically configurable automotive embedded systems The goals are to provide new platforms and methods that support scenarios such as software download (also during run-time), flexible internal resource configuration schemes (for availability or performance purposes), and flexible connectivity with external devices such as PDA's (for functionality and performance purposes). The main emphasis is on non-safety critical functions related to the telematics/infotainment domains, but in a connected Swedish national project the same scenarios are also considered for more safety critical functions. The challenge is to be able to fully exploit the flexibility of software while guaranteeing performance and dependability (including not distracting the driver).

Final Results

In many cases the technical achievements for Year 4 are continuations of work presented and described among the results of Year 3, i.e., in the previous section. Rather than repeating the same information we in those cases refer to the corresponding parts of the Year 3 section.

Achievement: Control of Server systems (LUND, UIUC, KTH)

The work on control of server systems in LUND has continued. Feedback in server systems has during last year's gained much interest in order to fulfil still increasing demands on performance and optimization regarding, for example, quality of service (QoS) requirements. In [1] Lund expand the feedback–based prediction scheme [2] for controlling a single server queue together with a new control strategy. These control structures have the benefit over other previously suggested control structures that no off-line estimation of the required work is needed. In addition, our solutions maintain or improve the performance, regarding average response time and loss of computational resources. At KTH work has been done on congestion control for small queues [14].

At UIUC Tarek Abdelzaher's group is working on feedback-based QoS control with applications to server systems. They are also developing new theory for real-time schedulability in open distributed systems based on aperiodic utilization bounds and synthetic utilization. The results are applicable to server systems under feedback control. The publications from UIUC are not reported here.

Achievement: Feedback-Based Resource Management in Cellular Devices (Ericsson, LUND)

Lund and Ericsson are continuing their collaboration in the national project "Feedback Based Resource Management and Code Generation for Soft Real-Time Systems". In the project a demonstrator based on dynamic resource control on a web-camera hosting multiple media streams is currently being developed. The support for reservation-based scheduling in the form of control groups within the new Linux kernel is also being investigated.

Achievement: Control and Optimization of Networked Systems (KTH, Ericsson, ABB, LUND)

KTH continues its activities on applying control and optimization to networked systems, and in particular wireless sensor and actuator networks. In this context they also study distributed optimization mechanisms for resource sharing. Part of this is performed in collaboration with ABB in the SOCRADES project and part of it with Ericsson and Swedish Institute for Computer Science (SICS). The work has led to the following publications [3-13, 15-20].

Achievement: Dynamically Configurable Automotive Embedded Systems (KTH, Volvo)

The work on dynamic reconfiguration of automotive embedded systems reported last year has continued and led to several joint publications. This has, however, been reported in the JPIA-Platform deliverable “Design Tools for Embedded Control” since it also relates to that activity.

Achievement: Resource Management

In this work UPVLC focus on dynamic memory management. Dynamic memory management can be considered in the proper way in real time applications due to the availability of dynamic memory allocators with constant allocation and deallocation. This work proposes a framework for handling dynamic memory in real-time systems. The framework provides both a flexible contract negotiation model to adapt the memory allocated to a set of task as closed as the required during the execution. Since UPVLC is doing this workly partly in collaboration with the University of York the publications resulting from this are reported in the network activity report Hard RT, ART and Control.

2.2.6 Testing and Verification Cluster

This cluster is composed of the following activities:

Testing and Verification Platform for Embedded Systems (Platform)
Led by Kim Larsen (Aalborg University / CISS)

Partner teams (leaders): Ed Brinksma (University of Twente), Pierre Wolper (Centre Fédéré de Verification), Philippe Schnoebelen (LSV), Thierry Jeron (INRIA), Yassine Lakhnech (Verimag), Wang Yi (Uppsala), Tom Henzinger (EPFL)

Affiliated Participants: Lubos Brim (University Brno), Henrik Leerberg (IAR Systems A/S), Tommy Ericsson (Telelogic), Jan Tretmans (Nijmegen), Sven H. Sørensen: (Motorola A/S), Thomas Hune: (Terma A/S)

Overview: The teams collaborating on this activity are leading tool providers for testing and verification, with particular emphasis on real-time, hybrid and stochastic aspects.

Automatic analysis of such quantitative aspects are crucial in validating embedded systems, but are computationally significantly more difficult than validation of simple functional aspects.

Thus, to address industrial size models continued development of new algorithmic techniques and data structures should be combined with powerful computational resources. We seek to establish this by maximal use, coordination and extension of existing local resources (e.g. PC-clusters) and by exploiting on-going work on exchange between and combinations of tools.

Despite advances in algorithmic techniques verification and test case generation are computationally notoriously hard problems.

Consideration of quantitative phenomena (real-time, stochastic) adds to the complexity. Thus, to address industrial size models powerful computational resources are necessary for example by maximal coordination of existing local resources.

The computational resources of the platform will initially be provided by existing powerful stand-alone computers with the various verification and testing tools being made available via a common web-based interface. A procedure for controlling access in a flexible and secure (e.g. in accordance with the individual tools licence agreements) manner will be investigated.

Among the tools that are candidates for beeing made available we mention: SPIN, SMV, UPPAAL, Kronos, Blast, TorX, TGV, FAST, CADP, IF; HyTech, visualSTATE, TAU, LASH, EMTCC and Rapture where the individual consortium member will have responsibility for integrating their tools into the platform.

The emerging advances in parallel and distributed model checking also motivate the development of a generally accessible server platform consisting of local clusters of (inexpensive) PCs.

Long term vision includes an experimental GRID infrastructure targeted specifically towards verification and testing.

Work in Year 1

During the first 12 months a number of improvements have been made on the individual tools as developed by the partners:

- The Vertecs team (IRISA) supports two test generation tools: TGV and STG. During the period, a new version of TGV (based on on-the-fly enumerative algorithms) linked to the IF toolbox (Verimag) has been developed using STL libraries (in place of CADP libraries).
- Results have been implemented in the TIMES tool for automated schedulability checking.
- CFV supports the verification tool LASH and hosts powerful servers dedicated to verification tools.
- A number of improvements have been made on the Uppaal real-time model checker (www.uppaal.com). This includes the possibility to enrich the timed automaton models with C code. An extension of Uppaal (Uppaal Cora), dedicated to solving optimal scheduling and planning problems, has been introduced. Recently, a version of Uppaal (Uppaal Tron), dedicated to online testing of real time systems, has been announced.

Also, a general distributed verification environment (DiVinE, Brno) has been deployed. The environment supports the development of distributed enumerative model checking algorithms, enables unified and credible comparison of these algorithms, and makes the distributed verification available for public use in a form of a distributed verification tool.

Finally, an overview of existing tools has been made accessible via a common web portal (the Yahooda web-page maintained by Brno).

Work in Year 2

Development of existing and new tools

Brno has completed deployment of the distributed verification tool *"DiVinE"* (version 0.7) for enumerative model checking of LTL properties on a network of workstations. This includes the development of new algorithms for cluster-based decomposition of state spac into strongly connected components to be used in reduction of state spaces

Neijmegen has recently implemented an initial extension of the *TorX* tool (*TorXakis*) for symbolic testing – based on the formalism of Symbolic Transition Systems.

IRISA has worked on symbolic test selection for extended automata using abstract interpretation and included the results by improving test selection in their toolset *STG*.

Verimag has continued work on conformance testing for real-time systems and in particular worked on general improvements on the tool *TTG* (Timed Test Generator).

A new version of *UPPAAL* (Aalborg, Uppsala), UPPAAL 4.0, has been released with a number of new facilities and algorithms *user defined functions* (syntax follows the style of C/C++/Java, and most control-flow constructs of C are supported), *priorities and channels* may be specified and dealt with during analysis, full support for *symmetry reduction* is implemented enabled by the introduction of a *scalar* datatype and the so-called *swep-line* method may be used to reduce memory consumption.

The online testing tool *Uppaal Tron* (Aalborg) has been ported to MS windows, and a new version 1.4 has been released. This represents a significant development effort since the OS and development environments on windows are quite different from those of Linux. We have identified specific technical problems with timing under windows. We believe that the windows version will greatly extend the applicability of the tool

A new variant of Uppaal, *Uppaal Tiga*, for the analysis and synthesis of winning strategies for times games has been released. Extensive evaluation of an experimental implementation of the algorithm yields very encouraging performance results.

Evaluation of tools

The planned work on tool dissemination and evaluation through case studies has been initiated through the establishment of an open repository for Artist2 Test and Verification Case Studies (https://bugsy.grid.aau.dk/artist2). The repository can be maintained by the individual tool providers and users through the use of Wiki.

Exploiting European Grid activities

ARTIST2 partners have participated in two European meetings on parallel and distributed model checking where the issue of exploiting grid activities to build a joint infrastructure has been discussed. The meetings showed that

- There are a number of ongoing European projects with respect to the usage of high performance and Grid-based servers for model checking. Each of the projects have made contributions through new distributed algorithms, new parallel architectures and new interesting applications, and it is likely that these activities will be their main focus for the immediate future. This means that the question of mutual exploitation of resources and the provision of a common web interface will be postponed for the time being.
- The long term vision of a joint high-performance verification platform is still relevant and should be maintained.
- There are already a number of facilities (e.g.) NorduGrid available that may be exploited be the individual tool providers and users. So far, a distributed version of Uppaal (DUppaal) has been made available on the NorduGrid in a certified manner via manual certificate distribution.

Work in Year 3

The technical results below are based on joint collaboration between the partners in terms of tool sharing and evaluation (all partners) and the development of new algorithms for test generation and controller synthesis (Twente, CFV, LSV, Aalborg). Also, IRISA, Brno, Aalborg and Uppsala are contributing with further tool development – partly based on the new algorithms., and Aalborg and Brno are investigating the high performance platform issue. Finally, the partners contribute to cross-cluster activities on tools and platform integration, and also interact with communities out ARTIST2.

Development of existing and new tools

IRISA has improved the STG tool (test generation for models with control and data) and it is now freely distributed (http://www.irisa.fr/vertecs/software.html#STG). Its integration with the NBAC analyzer and the APRON library has been improved. Also, a number of cases studies have been developped and experimented.

UPPSALA has developed a prototype tool (named CATS) for compositional timing and performance analysis using timed automata and the real time calculus developed at EPFL. It is based on an over-approximation technique in which a component of a system, modeled as a timed automaton is abstracted as a transducer of event streams described by arrival curves from the real-time calculus. This allows us to characterize the semantics of a system as a set of equations over streams. Many interesting properties such as schedulability and buffer

boundedness can be checked in solving the equations.The CATS tool is implemented in the Eclipse tool platform. As the main feature of the current version, it can be used to check the schedulability of a system and to estimate the best and worst case response times of its computation tasks. The tool is available for evaluation at www.timestool.com/cats.

AALBORG has continued its work on improving *UPPAAL* (www.uppaal.com) and its variants *UPPAAL Tron* (online testing) and *UPPAAL Tiga* (analysis and synthesis of winning strategies for timed games). In particular, the basic common DBM library has been improved and several industrial cases have been carried out. Also, the performance and the presentation of winning strategies have been improved significantly.

VERIMAG has implemented new techniques for deadlock detection in DeadlockFinder - a prototype tool that generates from BIP models sufficient conditions for deadlockfreedom.

BRNO has put a major effort into the development of tools to support parallel verification of complex embedded systems. The DiVinE tool has been made publically available and extended by a Promela front-end for SPIN compatible distributed model checking. With DiVinE Multi-Core (published at SPIN'07 and released in October 2007) the world-wide first parallel model-checker for multi-core architectures has been lunched. Furthermore, in the area of parallel techniques for the verification and analysis of embedded systems we focused on the development of new and enhanced algorithms for the enumerative parallel checking of reachibility, safety and liveness properties, in particular taking into account stochastic aspects of embedded systems.

Evaluation and dessimination of tools

A number of industrial case studies have been carried out by OFFIS, Twente, Aalborg and Uppsala. Apart from being documented via scientific papers, they are also collected and dessiminated through the open repository for Artist2 Test and Verification Case Studies (https://bugsy.grid.aau.dk/artist2). Links to mature versions of the applied tools may also be found at the web page.

Investigations wrt. High-performance tool server

In most verification cases, the performance bottle neck is the state space explosion, i.e. the size of the internal memory for representing the state space. As a result of questioning the development teams and experimenting with the tools, it turns out that most academic verification tools are limited by the fact that they have been developed for a 32-bit architecture. In fact, only a few of them are considering to become upgraded to 64-bit architectures, and our investigation only identified a single 64-bit tool, namely the SPIN model checker. The SPIN team is currently working on extending the tool to exploit a multi-cpu shared memory architecture.

Hence, the available tools for a possible high performance tool server are currently SPIN (on a shared memory architecture) and UPPAAL/DiVinE (on distributed PC clusters).

Aalborg is currently installing a large (600 node) PC cluster which will be made available via NoduGrid (www.nordugrid.org). It will mainly be dedicated to scientific computing, but experiments with high-performance model checking (also parallel and distributed model checking) will be possible to a certain extent.

Final Results

Development of existing and new tools

IRISA have improved the symbolic test generation tool, STG, and a new version can be downloaded from Inria Gforge: https://gforge.inria.fr/plugins/scmsvn/viewcvs.php/?root=bjeannet.

VERIMAG have applied their test generation tool TTG http://www-verimag.imag.fr/~krichen/ttg/index.html for the automatic generation of robotics observers. Also, they have demonstrated how transformations between different tool formats may be applied for verification of quantitative properties (BIP tool suite vs. FXML/Jahuel).

AALBORG have further developed the UPPAAL Tiga tool for controller synthesis. Especially, the tool now supports simulation of the identified winning strategi. Also, the UPPAAL real time model checker has been further improved and optimized, and a new version is being distributed in near future.

OFFIS have developed a modelling and verification framework for protocol validation and illustrated the concepts through a major industrial traffic communication protocol.

BRNO have further optimized the DiVinE tool and evaluated its performance and scalability on large-scale parallel systems. For preliminary experiments they have used the Distributed ASCI Supercomputer, a wide-area distributed system for Computer Science research in the Netherlands. As a next step they plan to make large scalability tests on the new cluster in Aaalborg. BRNO have extended the DiVinE tool with I/O efficient verification algorithms, which allow to exploit parallel hard disks. Quantitative LTL mode-checking has been added to the probabilistic version of DiVinE.

Evaluation and dessimination of tools

Industrial case studies have been carried out by OFFIS, TWENTE and Aalborg. Apart from being documented via scientific papers, they are also collected and dessiminated through the open repository for Artist2 Test and Verification Case Studies (https://bugsy.grid.aau.dk/artist2). Links to mature versions of the applied tools may also be found at the web page. In addition, Aalborg and OFFIS are involved in a number of industrial projects, where their own and also commercial tools are being applied in product development. E.g.:

- Uppaal Tron is being applied in the testing of new climate controllers for pig stables, and also in the testing of on-board satellite navigation software.
- Uppaal CORA (for cost optimal analysis) is being applied in the generation of cost optimal test cases for medical devices. As part of this development, a plugin for importing UML Stachart diagrams has been implemented.
- A commercial variant of UPPAAL (Vplus) is now being distributed for generating test cases to web services. The test cases provide plugins to commercial test execution tools like e.g. Mercury.

High Performance tool server

A common web interface for the tools supporting 64 bit architectures or distributed PC clusters has been developed and made accessible at https://benedict.grid.aau.dk/duppaal/ . In order to avoid misuse of the servers, one has to obtain certificate from the host organisation before it can be applied.

Quantitative Testing and Verification (NoE Integration)
Led by Ed Brinskma (University of Twente)

Partner teams (leaders): Kim G. Larsen (BRICS/Aalborg), Ed Brinksma (University of Twente), Pierre Wolper (Centre Fédéré de Verification), Philippe Schnoebelen (LSV), Thierry Jéron (INRIA/Rennes), Yassine Lakhnech (Verimag), Wang Yi (Uppsala), Tom Henzinger (EPFL), Werner Damm (OFFIS)

Affiliated Participants: Tretmans (Nijmegen), Bouajjani (LIAFA), Lubos Brim (University Brno), Tommy Ericsson (Telelogic), Sven H. Sørensen (Motorola A/S), Christer Nordstöm (ABB Automation), Jan Lindblad (Enea Embedded Technology), Alain Ourghanlian (EDF Recherche et Développement)

Overview: The long-term ambition of the Testing and Verification cluster is to improve current industrial practice for developing embedded systems applications by continuous dissemination and improvement of existing powerful testing and verification techniques. For embedded systems – besides functional correctness – properties concerning quantitative aspects including real-time constraints and constraints on quality of services are of utmost importance. It is therefore our aim to provide modelling formalisms, methods and tools which will allow such quantitative aspects to be dealt with at early design stages and utilized in a systematic (and ideally automatic) approach in the testing phase. Also, based on existing powerful (real-time) verification techniques new research challenges of industrial importance is taken-up including optimal scheduling, monitoring and fault diagnosis, coverage metrics, controller synthesis, analysis of hybrid models (allowing to take into account the physical environment in which an application is used) and robustness and implementability of timed models. The involved partners include leading European teams with responsibility for some of the most mature methods and tools for testing and verification of functional, timing and QoS properties.

There are several ongoing collaborations, including:

- The start of the STREP projects *Quasimodo* and *Multiform* mark significant collaboration between several partners of the cluster (Aalborg, Twente, CFV, LSV and Aalborg, Twente, Verimag respectively). Also a number of teams affiliated with the cluster are partners in the proposal (Aachen, Saarlandes, ESI, Nijmegen).
- Numerous collaborations between LSV and Verimag on national projects (Eva, Prouvé, Rossignol, Action Spécifique du CNRS).
- Aalborg and Uppsala has since 95 continuously collaborated on the development of the tool UPPAAL in parallel with the development of Kronos at Verimag. In particular the collaboration has lead to a spin-off company (officially named UP4ALL).
- LSV, Aalborg and Twente are collaborating on problems related to optimal control and scheduling for real-time systems.
- CVF and Aalborg are collaborating on controller synthesis for real-time systems under partial observability and for efficient realization of synthesis algorithms within the verification tool UPPAAL.
- CVF, LSV and Aalborg are partners in the newly started ESF project GASICS: Games for Analysis and Synthesis of Interactive Computational Systems.
- Twente and INRIA have long been collaborating on testing methodologies and tools;
- INRIA and Verimag has for a long time collaborated on developing the testing tool TGV, and are currently collaborating on connecting IF and TGV within the Agedis IST project and the national project AS Testic
- CVF and EPFL have numerous collaborations on controller synthesis as well as analysis of stochastic model.

Work in Year 1

Work carried out in the first months include:

- The Vertecs team of INRIA has worked on test generation for models of infinite state systems with control and data. Systems are modelled with ioSTS (e.g. automata extended with data). Test generation from specification models and test purposes is based on syntactic transformations guided by approximate co-reachability analysis. The main achievements has been a new formalisation of symbolic test generation and a combination of verification and testing for safety properties.
- Uppsala has shown that the schedulability problem will be undecidable if tasks execution times may vary within an interval (representing the best and worst case execution times). They also developed an algorithm to compute the worst-case response times of non-uniformly recurring fixed-priority tasks. For systems containing only periodic tasks, the algorithm performs as well as the classic method for Rate-Monotonic Analysis. These results have been implemented in the TIMES tool for automated schedulability checking.
- A number of improvements have been made on the UPPAAL real-time model checker (www.uppaal.com). This includes the possibility to enrich the timed automaton models with C code. (Aalborg) This has given an important increase in the expressiveness of the modelling tool, e.g. the possibility to include advanced data types. During the period, the tool has been applied for off-line test generation on a connectivity testing framework.
- An extension of UPPAAL (UPPAAL Cora, Aalborg), dedicated to solving optimal scheduling and planning problems, has been introduced. This version is based on a version of the classical timed automaton formalism extended with auxiliary cost variables and with a modified version of the UPPAAL verification engine to take the accumulation of cost into account. During the period, several new algorithms have been designed for transforming the cost optimisation problem into a max-flow problem (in stead of a linear programming problem), and they will be introduced in forthcoming versions of the tool.
- Twente has carried out work on
    - Scheduling by reachability analysis: The feasibility of using search techniques from model checking to synthesize and analyse scheduling problems of industrial relevance was established.
    - Integrated quantitative analysis: The usefulness of model checking techniques for Markov chain analysis was further extended by application to Markov reward modelling. An industrial case study was carried out concerning an availability monitoring algorithm for self-configuring networks, with analysis carried out using the MODEST modelling formalism and the Moebius tool.
    - Modelling of hybrid systems: A process algebraic formalism for the modelling and analysis of hybrid systems has been developed.
    - Real-time testing: A real-time testing theory for quiescent systems has been formulated, implemented as a TorX extension, and extended to multi-channel interfaces.
- Information on formal methods relevant for industrial applications have been collected by OFFIS, and support was given to industrial partners to perform case studies on formal verification tools (commercial ones as well as academic ones). The work on case studies showed that it actually is possible to formally prove safety properties of e.g. existing car steering control software.

- Uppsala has developed a sampling semantics for timed automata, and shown that the new semantics gives rise to a natural notion of digitalization for timed lnguages. A recent result shows that the language inclusion problem in this setting is decidable, which in turn implies that for any timed automaton, a digital machine can be constructed systematically, which accepts the digitalized language of the automaton.
- A version of UPPAAL (UPPAAL Tron, Aalborg), dedicated to online testing of real time systems, has been announced. By using UPPAAL Tron, one can extend the testing power of traditional tools substantially, partly because one can run tests for a very long time, and also because Uppaal Tron gives the possibility to build various stochastic criteria into the test selection algorithm. During the period, further performance improvements have been made, and also a first realistic industrial case study has been made. The purpose of the study was to test the functionality of an existing electronic cooling thermostate, and several inconsistencies wrt. the product specification were revealed.
- Cachan has designed techniques for computing the convex hull of Presburger-definable sets of tuples of integers. These abstraction techniques are used in model-checking of complex counter systems; Improved techniques for verification of communicating systems including half-duplex channel systems and probabilistic lossy channel systems; Introduced the concept of "flat acceleration", a powerful generic algorithmic approach for the symbolic computation of reachability sets in regular model checking; In-depth study of the descriptive power of formalisms based on timed-automata and extensions, constrasted with verification costs; Model checking sets of paths: an approach that sits in between test and model checking. Also, quantitative analysis of priced timed automata, and used timed automata as a tool for fault diagnosis; Designed new probabilistic models supporting improved verification algorithms; Extensions of temporal logic formalisms, and associated verification techniques; Used UPPAAL for the verification of a multitask automation system.

Work in Year 2

The following is a short summary of work carried out in the second year:

- We have released UPPAAL 4.0 being the result of over two and half years of development and contains many new features, additions to the modelling language, performance improvements, enhancements and polish to the easy to use graphical user interface, and libraries are available free of charge for academic, educational and evaluation purposes
- We have studied channel systems whose behaviour (sending and receiving messages via unbounded FIFO channels) must follow given timing constraints specifying the execution speeds of the local components.
- We have presented an algorithm for inferring a timed-automaton model of a system from information obtained by observing its external behavior. In this work, the full class of event-recording automata has been considered.
- We have worked on symbolic test selection for extended automata using abstract interpretation.
- We have worked on symbolic Determinisation of Extended Automata.
- We have implemented tool support for off-line test generation for real-time systems in UPPAAL Cover and UPPAAL Tron and applied it to a case study where a model-based approach to black-box testing is applied to verify that a Wireless Application Protocol (WAP) gateway conforms to its specification. The WAP gateway is developed by

Ericsson and used in mobile telephone networks to connect mobile phones with the Internet.

- We have worked on conformance testing of programs with floating point numbers with respect to its specification with real numbers.
- We have worked on black-box testing of cryptographic protocols, using a compositional approach for checking secrecy and authenticity properties of cryptographic protocols integrating ideas from verification, conformance testing, and learning, with application to biometric passports.
- Work on verification of communication protocols using abstractiInterpretation of FIFO queues.
- Work on supervisory control of infinite symbolic systems using abstract interpretation.
- We have worked on analysis of priced (Weighted) timed automata, in particular settling decidability of optimal reachability in presence of multi cost functions and proved undecidability of model checking and optimal control in general (for priced timed automata with 3 or more clocks)
- We have worked on efficient implementation of cost-optimal reachability for priced timed automata using a symbolic A* algorithm in UPPAAL Cora.
- Work on robustness issues for timed and hybrid automata by introduction of a parametric semantics for timed controllers called the ASAP semantics.
- We have worked on analysis of O-minimal Hybrid Systems, refinement of abstraction for affine hybrid automata and development of an acceleration method suited for linear hybrid automata.
- We have developed and implemented an efficient on-the-fly algorithm for solving timed games wrt reachability and safety propertis. The implementation is available in the tool UPPAAL Tiga.
- For finite games we have proposed a fixed point theory of anti-chains to efficiently solve games of imperfect information.
- We have proposed algorithms for the verification of infinite state systems including rectangular abstrqctions of hybrid automata.
- We haved defined Quantitative similarity between timed systems and proposed logics for real-time games allowing to specify objectives.
- We have defined the formalism of Symbolic Transition Systems (STS) in order to support testing of systems with data. Also to support testing of communication protocls a testing theory allowing for action refinement was proposed.
- We have developed a framework for test coverage semantics as well as a testing theory for probabilistic processes.
- We have developed on-line testing of real-time systems in the tool UPPAAL Tron as well as a theory for conformance testing for real-time systems as used in the tool TTG.
- We have developed a framework for compositional reasoning about qualitatitve system properties.
- We have proposed a symbolic algorithm for the analysis of the robustness of timed automata, that is the correctness of the model in presence of small drifts on the clocks or imprecision in testing guards.

- We have presented a novel approach to synthesize good schedules for a class of scheduling problems that is slightly more general than certain existing scheduling problems.
- We have presented an (semi-decision procedure) algorithm for cost-bounded probabilistic reachability problem.
- We have studied the state identification problems for finite-state transducers, and the fundamental observation problem of decentralized observation.'
- We have provided an automatic method for calculating the path condition for programs with real time constraints. This method can be used for the semiautomatic verification of a unit of code in isolation, i.e., without providing the exact values of parameters with which it is called.
- We have showed how Allen's logic can be translated to LTL and how to synthesize automatically monitors for specifications in this logic.
- We have developed a framework for development and validation of product lines. In the approach families of embedded discrete finite state programs are modeled using input-enabled alternating transition systems. One model describes all functionality, while each variant is defined by an environment, describing its possible uses.
- We have worked on compositional verification using I/O-Automata

Work in Year 3

The following is a short summary of what was carried out during the third year:

- We implemented efficient on-the-flh algorithms for solving timed games with respect to reachability and safety properties. The implementation has resulted in the branch UPPAAL TIGA providing a mature and fully integrated tool capable of synthesizing winning strategies for models exploiting the full modeling language of UPPAAL 4.0. Strategies can be represented as BDDs and CDDs providing compact formats for control code.
- An industrial application of UPPAAL TIGA to the synthesis of a climate controller for pig stables have been conducted. The resulting tool chain (UPPAAL TIGA and Simulink) has been applied by the company Skov A/S.
- We have shown decidability, provided an efficient on-the-fly algorithm for synthesizing stategies for timed games with partial observability. Also applications have been provided. The algorithm will be integrated with UPPAAL TIGA by the end of 2008.
- We have shown how to reduce various simulation and alternating simulation preorders between timed automata and timed game automata to safety games.
- The game-theretic approach has been applied to the testing of uncontrollable real-time systems. Specifying test purposes as TCTL fomulas, UPPAAL TIGA has been exployided to synthesis test strategies. Case studies have been conducted.
- Lower and upper bound complexity results for refinement and consistency of modal transition systems have been obtained.
- Model checking and optimal reachability for one-clock priced timed automata and priced tmed games have been shown decidable.

- Improved state space search algorithms for timed automata models – exploiting heuristic search and agent based techniques – have been described and implemented within the UPPAAL verification engine.
- The DBM Library of UPPAAL – which provided efficient implementation of the symbolic datastructure used for exploring timed automata state spaces – has been significantly improved.
- We have given formal semantics for dynamic fault trees (DFT) and developed compositional analysis techniques aliviating the state space explosion problem.
- We have developed equivalence relations and metrics for concurrent, stochastic games.
- We have proposed a symbolic algorithm for the analysis of the robustness of timed automata. Prototype implementation within UPPAAL has been made.
- We developed an semi-decision algorithm for cost-bounded probabilistic reachability in timed automata extended with prices (on edges and locations) and discrete probabilistic branching.
- We developed a framework for quantitative reasoning about quantitative systems. More precisely, we developed quantitative logics and quantitative refinement relations and showed their connections.
- A temporal interface for a component is a finite automaton that specifies the legal sequences of calls to functions that are provided by the component. We compared and evaluated three different algorithms for automatically extracting temporal interfaces from code
- We provided efficient methods for synthezing control for reactive systems by solving two-player games on graphs. The efficiency of the method avoids expensive determinization of (Büchi) automata.
- We have extended the game logic ATL with time quantifiers allowing objectives on real-time games to be specified.
- We studied stochastic graph games with omega-regular winning conditions specified as Rabin or Streett objectives. We developed a compositional theory of system verification, where specifications assign real-numbered costs to systems.
- We presented a verification methodology for cooperating traffic agents covering analysis of cooperation strategies, realization of strategies through control, and implementation of control.
- We analysed a flap controller (high-lift) case study. This application is derived from a case study for Airbus, a controller for the flaps of an aircraft.
- We analysed a distributed controller for Tramways. The "Single-tracked Line Segment" (SLS) case study stems from an industrial partner of the UniForM-project.
- The STG tool (test generation for models with control and data) has been improved, and a number of cases studies have been developped and experimented.
- A methodology integrating verification and conformance testing for the formal validation of reactive systems has been described.
- We have proposed an extension of the model of communicating automata (CFSM): Symbolic Communicating Machines (SCM), where messages carry data in infinite

domains, and an approximate reachability analysis method on this model, based on lattice automata.

- Approximation and abstraction methods for the validation of for timed systems with respect to particular resource-related issues covering a broad range of resources such as processors, buffers and memory blocks etc.
- Ideas underlying our new algorithms for controller synthesis under imperfect information has recently been extended to solve classical problems in automata theory for finite word languages and infinite words.
- We have defined an new abstract fixpoint checking algorithm with automatic refinement by backward completion in Moore closed abstract domains. We have studied the properties of our algorithm and prove it to be more precise than the counterexample guided abstract refinement algorithm (CEGAR).
- Development of an acceleration method suited for the analysis of linear hybrid automata.
- Testing distributed systems through symbolic model checking
- Study of the properties of automata-based representation of sets of real numbers.
- A new technique for computing the convex hull of an automaton-represented finite set of integers was introduced.
- Development of a method for test case generation for ultimately periodic paths.
- Development of a conformance testing framework for hybrid systems, including definition of coverage measure for hybrid systems and algorithms for coverage-guided test generation.
- We developed new composability and compositionality techniques for deadlock-freedom implemented within BIP.
- Developed an abstract domain extending differerence-bounded matrices with disequality constraints.

Final Results

-- Aalborg --

A Game-Theoretic Approach to Real-Time System Testing

This work presents a game-theoretic approach to the testing of uncontrollable real-time systems. By modelling the systems with Timed I/O Game Automata and specifying the test purposes as Timed CTL formulas, we employ a recently developed timed game solver UPPAAL-TIGA to synthesize winning strategies, and then use these strategies to conduct black-box conformance testing of the systems. The testing process is proved to be sound and complete with respect to the given test purposes. Case study and preliminary experimental results indicate that this is a viable approach to uncontrollable timed system testing.

Infinite Runs in Weighted Timed Automata with Energy Constraints
with LSV, ENS Cachan, France

We study the problems of existence and construction of infinite schedules for finite weighted automata and one-clock weighted timed automata, subject to boundary constraints on the accumulated weight. More specifically, we consider automata equipped with positive and

negative weights on transitions and locations, corresponding to the production and consumption of some resource (e.g. energy). We ask the question whether there exists an infinite path for which the accumulated weight for any finite prefix satisfies certain constraints (e.g. remains between 0 and some given upper-bound). We also consider a game version of the above, where certain transitions may be uncontrollable.

Complexity of Decision Problems for Mixed and Modal Specifications
with ITU, Copenhagen and Imperial College London, UK

We consider decision problems for modal and mixed transition systems used as specifications: the common implementation problem (whether a set of specifications has a common implementation), the consistency problem (whether a single specification has an implementation), and the thorough refinement problem (whether all implementations of one specification are also implementations of another one). Common implementation and thorough refinement are shown to be PSPACE-hard for modal, and so also for mixed, specifications. Consistency is PSPACE-hard for mixed, while trivial for modal specifications. We also supply upper bounds suggesting strong links between these problems.

Testing Real-Time Systems Using UPPAAL
with Uppsala University

This chapter presents principles and techniques for model-based black-box conformance testing of real-time systems using the Uppaal model-checking tool-suite. The basis for testing is given as a network of concurrent timed automata specified by the test engineer. Relativized input/output conformance serves as the notion of implementation correctness, essentially timed trace inclusion taking environment assumptions into account. Test cases can be generated offline and later executed, or they can be generated and executed online. For both approaches this chapter discusses how to specify test objectives, derive test sequences, apply these to the system under test, and assign a verdict.

Fast Directed Model Checking Via Russian Doll Abstraction,
with University of Friburg and University of Innsbruck

Directed model checking aims at speeding up the search for bugs in a system through the use of heuristic functions. Such a function maps states to integers, estimating the state's distance to the nearest error state. The search gives a preference to states with lower estimates. The key issue is how to generate good heuristic functions, i.e., functions that guide the search quickly to an error state. An arsenal of heuristic functions has been developed in recent years. Significant progress was made, but many problems still prove to be notoriously hard. In particular, a body of work describes heuristic functions for model checking timed automata in Uppaal, and tested them on a certain set of benchmarks. Into this arsenal we add another heuristic function. With previous heuristics, for the largest of the benchmarks it was only just possible to find some (unnecessarily long) error path. With the new heuristic, we can find provably shortest error paths for these benchmarks in a matter of seconds. The heuristic function is based on a kind of Russian Doll principle, where the heuristic for a given problem arises through using Uppaal itself for the complete exploration of a simplified instance of the same problem. The simplification consists in removing those parts from the problem that are distant from the error property. As our empirical results confirm, this simplification often preserves the characteristic structure leading to the error.

Model-checking one-clock priced timed automata,
with LSV, ENS Cachan, France

We consider the model of priced (a.k.a. weighted) timed automata, an extension of timed automata with cost information on both locations and transitions, and we study various model-checking problems for that model based on extensions of classical temporal logics with cost constraints on modalities. We prove that, under the assumption that the model has only one clock, model-checking this class of models against the logic WCTL, CTL with cost-constrained modalities, is PSPACE-complete (while it has been shown undecidable as soon as the model has three clocks). We also prove that model-checking WMTL, LTL with cost-constrained modalities, is decidable only if there is a single clock in the model and a single stopwatch cost variable (i.e., whose slopes lie in {0,1}).

Optimal infinite scheduling for multi-priced timed automata,
with LSV, ENS Cachan, France and ESI, Eindhoven NL

This paper is concerned with the derivation of infinite schedules for timed automata that are in some sense optimal. To cover a wide class of optimality criteria we start out by introducing an extension of the (priced) timed automata model that includes both costs and rewards as separate modelling features. A precise definition is then given of what constitutes optimal infinite behaviours for this class of models. We subsequently show that the derivation of optimal non-terminating schedules for such double-priced timed automata is computable. This is done by a reduction of the problem to the determination of optimal mean-cycles in finite graphs with weighted edges. This reduction is obtained by introducing the so-called corner-point abstraction, a powerful abstraction technique of which we show that it preserves optimal schedules.

In this paper, we prove the decidability of the minimal and maximal reachability problems for multi-priced timed automata, an extension of timed automata with multiple cost variables evolving according to given rates for each location. More precisely, we consider the problems of synthesizing the minimal and maximal costs of reaching a given target location. These problems generalize conditional optimal reachability, i.e., the problem of minimizing one primary cost under individual upper bound constraints on the remaining, secondary, costs, and the problem of maximizing the primary cost under individual lower bound constraints on the secondary costs. Furthermore, under the liveness constraint that all traces eventually reach the goal location, we can synthesize all costs combinations that can reach the goal.

Optimal reachability for multi-priced timed automata

The decidability of the minimal reachability problem is proven by constructing a zone-based algorithm that always terminates while synthesizing the optimal cost tuples. For the corresponding maximization problem, we construct two zone-based algorithms, one with and one without the above liveness constraint. All algorithms are presented in the setting of two cost variables and then lifted to an arbitrary number of cost variables.

Development of UPPAAL

In 2008 the concrete simulator of UPPAAL-TIGA was improved. It is now more stable and its interface has been updated. Another completely new algorithm was implemented to handle timed games with partial observability. It is now possible to define actions on edges inside the graphical editor and define observations when checking properties. The implementation only need to have loop detections added to be complete w.r.t. our paper. A second new major feature was also the implementation of timed games with Buchi accepting states, while avoiding some other bad states. This new algorithm is on-the-fly in the sense that it can stop whenever it has found a stable set of accepting and winning states. It works in two stages

where a fix-point on the set of winning states is computed and then a fix-point on the set of winning states that are Buchi accepting.

Concerning UPPAAL, the engine is now able to merge DBMs dynamically when exploring the state-space. This is a transparent feature for the user. This is done automatically whenever possible.

Another new feature has been the addition of stop-watches. It is now possible to add to locations expressions of the form "x'==expr" where x is a clock and expr an expression evaluating to 0 or 1. There is no other needed syntax additions. Any clock can be stopped. However, the algorithm used becomes an over-approximation whenever a state that is stopping a clock is reached.

We also mention that the DBM library has been updated internally to cope with the extensions we have made. A new version will be released soon.

Discount-Optimal Infinite Runs in Priced Timed Automata

Discount-optimal infinites scheduling for priced timed automata has been shown decidable using region-based techniques (so-called corner-point abstraction). Using discounting in optimization criteria is a often used in Control Theory, and leads to a simple fixed-point characterization in the setting of weighted timed automata. The fixed-point characterization suggests an efficient algorithm in contrast to limit-ratio optimality.

Off-line test case generation

Off-line test-case generation from I/O timed automata models using increasingly techniques depending on properties of the given model. The techniques ranges from model checking (suitable if the model is controllable, i.e. deterministic and has neither timing uncertainty nor conflicting outputs), synthesis of testing strategy (suitable if the model is deterministic but fully observable), synthesis of strategy under partial observability. Also, testing strategies with respect to a given test purpose but relying on corporation from the system under test have been given.

Timed Interface Theory

An interface theory for real-time systems using timed games have beedn developed. Here UPPAAL TIGA supports compatibility, refinement as well as consistency checking of component specifications.

Slicing for UPPAAL.

The focus of this thesis is to introduce slicing for Uppaal. Slicing is a technique based on static analysis used to reduce the syntactic size of models or applications. In this thesis, we show how slicing may be used to construct reachability preserving reductions of Uppaal models possibly improving the performance of the tool. Using automated slicing in Uppaal will eliminate the need for users to manually optimize models for faster verification of a certain property. Moreover, it allows less experienced users of Uppaal, which unknowingly may design models, containing unnecessary large amounts of data, to verify properties which \uppaal otherwise would have been unable to check.

Design Verification Patterns

Design Verification Patterns are formal specifications that define the semantics of design patterns. For each design pattern, the corresponding verification pattern give a set of proof

obligations. They must be discharged for a correct implementation of the pattern. Additionally there is a set of properties that may be used in the design and verification of applications that employ the pattern. The concept is illustrated by examples from general software engineering and more specialised properties for embedded software.

Model-based Schedulability Analysis of Safety Critical Hard Real-Time Java Programs

This report describes the implementation of SARTS, a model based schedulability analysis tool used for hard real-time systems. SARTS is used to translate hard real-time systems, implemented in Java, to a timed automata model in UPPAAL.

The system being analyzed must be implemented in SCJ2, a safety critical profile for Java developed in this project, based on SCJ. The target hardware is the time predictable Java processor JOP, developed specifically for hard real-time systems.

Several experiments have been conducted to illustrate the accuracy of SARTS compared to existing tools. It is shown how the model based approach can result in a more accurate analysis, than possible with traditional analyses.

-- Twente --

Quantative reasoning frameworks

We have further refined and extended our quantitative verification framework. This framework allows us to reason about quantitative properties in a quantitative way, that is, rather than saying if a property holds for a system, we express to what extend the property holds for a system We have developed a quantitative logic QLTL, which is a quantitative counterpart of QLTL; and in we have further developed our work on metrics for games and on the logic QLTL. Moreover, we have used this framework to define a testing theory that describes how to test systems in the presence of measurement imprecisions.

Architectural dependability analysis

We have developed an architectural language that allows one to reason about dependability properties at an architectural level. By annotating an existing system design with dependability information (such as failure rater for components, or recovery times for processes), our method allows the designer to extract automatically analytical models from which various dependability measures (such as availability) can be computed automatically.

Testing Probabilistic Processes

We have defined a testing scenario that characterized in when two probabilistic processes exhibit the same visible behavior, i.e. when these are trace distribution equivalent. Along the way, we establish an Induction-Approximation theorem, which is of independent interest, since it related properties of infinite executions to properties of finite executions.

-- CVF --

Synthesis with incomplete information
(cooperation with EPFL, U Aalborg, and EC Nantes)

We have continued our collaboration with EPFL on algorithms for the synthesis of controller with imperfect information. In this research, we have proposed a fixed point theory to solve games of imperfect information. Â The fixed point theory is defined on the lattice of antichains

of sets of states. Â Contrary to the classical solution proposed by Reif, our new solution does not involve determinization. Â As a consequence, it is readily applicable to classes of systems that do not admit determinization. Notable examples of such systems are timed and hybrid automata. Â As an application, we show that the discrete control problem for games of imperfect information defined by rectangular automata is decidable. This result extends a result by Henzinger and Kopke.

Those results have been extended to stuttering invariant and observation based strategy in collaboration with U Aalborg and EC Nantes. These results should be integrated into the tool UppAal-Tiga in 2008.

Improved algorithms for the automata-based approach to model-checking
(in collaboration with EPFL)

Ideas underlying our new algorithms for controller synthesis under imperfect information have recently been extended to solve classical problems in automata theory for finite word languages and infinite words. With this new method, inclusion between two nondeterministic automata can be solved much more efficiently than with previously known algorithms. Those results should lead to the development of a new model-checking tool for linear time specifications expressed in LTL or using nondeterministic Buechi automata.

Fixed point based abstraction refinement
(in collaboration with ENS Paris)

We have defined an new abstract fixpoint checking algorithm with automatic refinement by backward completion in Moore closed abstract domains. Â We have studied the properties of our algorithm and prove it to be more precise than the counterexample guided abstract refinement algorithm (CEGAR). Â Contrary to several works in the literature, our algorithm does not require the abstract domains to be partitions of the state space. We have shown that our automatic refinement technique is compatible with so-called acceleration techniques. Â Furthermore, the use of Boolean closed domains does not improve the precision of our algorithm.

Development of an acceleration method suited for linear hybrid automata.

This method generalizes previous work on acceleration of integer-based systems, and provides a semi-algorithm for exploring the Â state-space of general linear hybrid automata, without abstracting away parts of the system or performing approximations. This method has been shown to be complete over the specific subclass of timed automata, but is also applicable to a much broader class of systems.

New efficient approximate verification based on symmetry markers.

This new verification technique can be used in various model-checker and in particular the spin tool. Â It exploits state-space symmetries induced by scalarset values used in a model. The technique involves efficiently computing a marker for each state encountered during search. A complete verification method only partially exploits symmetry; an approximate verification method fully exploits symmetry. We describe how symmetry markers can be efficiently computed and integrated into the SPIN tool. Â An empirical evaluation of our technique shows very good performance results and a high degree of precision for the approximate method (i.e. very few non-symmetric states receive the same marker). We also identify a class of models for which the approximate technique is precise.

Testing Distributed Systems through Symbolic Model Checking

The observation of a distributed system´s finite execution can be abstracted as partial order trace. We show that testing that such a distributed execution satisfies some global property amounts therefore to model check the corresponding trace. We provide an efficient symbolic CTL model checking algorithm for traces. This method is based on a symbolic data structure, called Interval Sharing Trees, allowing efficiently representing and manipulating sets of k-uples of naturals. Efficient symbolic operations are defined on this data structure in order to deal with all CTL modalities.

Study of the properties of automata-based representations of sets of real numbers

Automata-based representations of sets of real vectors are useful for manipulating the sets of configurations of infinite-state systems during state-space exploration. We have established that the sets of real vectors that can be represented by weak deterministic automata in all integer bases are exactly those thatare definable in first-order additive arithmetic. This generalizes to real numbers the well-known Cobham's theorem on the representability of sets of integers, and provides a theoretical justification to the use of weak deterministic automata as data structures for representing sets of reals in actual verification applications.

Alternative semantics for timed automata.

Like most models used in model-checking, timed automata are an idealized mathematical model used for representing systems with strong timing requirements. In such mathematical models, properties can be violated, due to unlikely (sequences of) events. We propose two new semantics for the satisfaction of omega-regular properties, one based on probabilities, and the other one based on topology, to rule out these sequences. We prove that the two semantics are equivalent and lead to a PSPACE-Complete qualitative model-checking problem for LTL over finite executions. We also obtain decidability results for both qualitative and quantitative problems on infinite runs for one-clock timed automata.

Model-checking TATL on TCGS.

We propose a new model for timed games, based on concurrent game structures (CGSs). Compared to the classical timed game automata Â of Asarin et al., our timed CGSs are "more concurrent", in the sense that they always allow all the agents to act on the system, independently of the delay they want to elapse before their action. Timed CGSs weaken the "element of surprise" of timed game automata reported by de Alfaro et al. We prove that our model has nice properties, in particular that model-checking timed CGSs against timed ATL is decidable via region abstraction, and in particular that strategies are "region-stable" if winning objectives are. We also propose a new extension of TATL, containing ATL*, which we call TALTL. We prove that model-checking this logic remains decidable on timed CGSs. Last, we explain how our algorithms can be adapted in order to rule out Zeno (co-)strategies, based on the ideas of Henzinger et al.

Study of the properties of automata-based representations of sets of real numbers

Automata-based representations of sets of real vectors are useful for manipulating the sets of configurations of infinite-state systems during state-space exploration. We have established that the sets of real vectors that can be represented by infinite-word automata in all integer bases are exactly those that are definable in first-order additive arithmetic. As an important corollary, such sets can also be represented by weak deterministic automata, which can be used as actual data structures in verification applications. This result generalizes to real

numbers the well-known Cobham's theorem on the representability of sets of integers, and provides a theoretical justification to the use of weak deterministic automata.

-- EPFL --

Timed parity games: Complexity and robustness

We considered two-player games played in real time on game structures with clocks and parity objectives. The games are concurrent in that at each turn, both players independently propose a time delay and an action, and the action with the shorter delay is chosen. To prevent a player from winning by blocking time, we restricted each player to strategies that ensure that the player cannot be responsible for causing a zeno run. First, we presented an efficient reduction of these games to turn-based (i.e., nonconcurrent) finite-state (i.e., untimed) parity games. The states of the resulting game are pairs of clock regions of the original game. Our reduction improved the best known complexity for solving timed parity games. Moreover, the rich class of algorithms for classical parity games can now be applied to timed parity games. Second, we considered two restricted classes of strategies for the player that represents the controller in a real-time synthesis problem, namely, limit-robust and bounded-robust strategies. Using a limit-robust strategy, the controller cannot choose an exact real-valued time delay but must allow for some nonzero jitter in each of its actions. If there is a given lower bound on the jitter, then the strategy is bounded-robust. We showed that exact strategies are more powerful than limit-robust strategies, which are more powerful than bounded-robust strategies for any bound. For both kinds of robust strategies, we presented efficient reductions to standard timed automaton games. These reductions provide algorithms for the synthesis of robust real-time controllers.

Trading infinite memory for uniform randomness in timed games

We considered concurrent two-player timed automaton games with omega-regular objectives specified as parity conditions. These games offer an appropriate model for the synthesis of real-time controllers. Earlier works on timed games focused on pure strategies for each player. We studied, for the first time, the use of randomized strategies in such games. While pure (i.e., nonrandomized) strategies in timed games require infinite memory for winning even with respect to reachability objectives, we showed that randomized strategies can win with finite memory with respect to all parity objectives. Also, the synthesized randomized real-time controllers are much simpler in structure than the corresponding pure controllers, and therefore´easier to implement. For safety objectives we proved the existence of pure finite-memory winning strategies. Finally, while randomization helps in simplifying the strategies required for winning timed parity games, we proved that randomization does not help in winning at more states.

Minimum-time reachability in timed games,
EPFL and CVF

We considered the minimum-time reachability problem in concurrent two-player timed automaton game structures. We showed how to compute the minimum time needed by a player to reach a target location against all possible choices of the opponent. We did not put any syntactic restriction on the game structure, nor did we require any player to guarantee time divergence. We only required players to use receptive strategies which do not block time. The minimal time is computed in part using a fixpoint expression, which we showed can be evaluated on equivalence classes of a non-trivial extension of the clock-region equivalence relation for timed automata.

Quantitative languages, EPFL:Quantitative generalizations of classical languages, which assign to each word a real number instead of a boolean value, have applications in modeling resource-constrained computation. We used weighted automata (finite automata with transition weights) to define several natural classes of quantitative languages over finite and infinite words; in particular, the real value of an infinite run is computed as the maximum, limsup, liminf, limit average, or discounted sum of the transition weights. We defined the classical decision problems of automata theory (emptiness, universality, language inclusion, and language equivalence) in the quantitative setting and study their computational complexity. As the decidability of language inclusion remains open for some classes of weighted automata, we introduced a notion of quantitative simulation that is decidable and implies language inclusion. We also gave a complete characterization of the expressive power of the various classes of weighted automata. In particular, we showed that most classes of weighted automata cannot be determinized.

Controller synthesis with budget constraints

We studied the controller synthesis problem under budget constraints. In this problem, there is a cost associated with making an observation, and a controller can make only a limited number of observations in each round so that the total cost of the observations does not exceed a given fixed budget. The controller must ensure some omega-regular requirement subject to the budget constraint. Such budget constraints arise in designing and implementing controllers for resource-constrained embedded systems, where a controller may not have enough power, time, or bandwidth to obtain data from all sensors in each round. Budget constraints lead to games of imperfect information, where the unknown information is not fixed a priori, but can vary from round to round, based on the choices made by the controller how to allocate its budget. We showed that the budget-constrained synthesis problem for omega-regular objectives is complete for exponential time. In addition to studying synthesis under a fixed budget constraint, we studied the budget optimization problem, where given a plant, an objective, and observation costs, we have to find a controller that achieves the objective with minimal accumulated cost (or minimal peak cost). We showed that this problem is reducible to a game of imperfect information where the winning objective is a conjunction of an omega-regular condition and a long-run average condition (or a least max-cost condition), and this again leads to an exponential-time algorithm. Finally, we extended our results to games over infinite state spaces, and show that the problems are decidable for infinite state games with stable quotients of finite index. Consequently, the discrete time budget-constrained synthesis and budget optimization problems are decidable for rectangular hybrid automata.

Model checking omega-regular properties of interval Markov chains

We studied the problem of model checking Interval-valued Discrete-time Markov Chains (IDTMC). IDTMCs are discrete-time finite Markov Chains for which the exact transition probabilities are not known. Instead in IDTMCs, each transition is associated with an interval in which the actual transition probability must lie. We considered two semantic interpretations for the uncertainty in the transition probabilities of an IDTMC. In the first interpretation, an IDTMC represents a (possibly uncountable) family of (classical) discrete-time Markov Chains, where each member of the family is a Markov Chain whose transition probabilities lie within the interval range given in the IDTMC. We call this semantic interpretation Uncertain Markov Chains (UMC). In the second semantics for an IDTMC, which we call Interval Markov Decision Process (IMDP), the uncertainty is resolved through non-determinism. In other words, each time a state is visited, we adversarially pick a transition distribution that respects the interval constraints, and take a probabilistic step according to the chosen distribution. We introduced a logic, omega-PCTL, that can express liveness, strong fairness, and omega-regular properties (such properties cannot be expressed in PCTL). We showed that the omega-PCTL model

checking problem for Uncertain Markov Chain semantics is decidable in PSPACE (same as the best known upper bound for PCTL), and for Interval Markov Decision Process semantics it is decidable in coNP (improving the previous known PSPACE bound for PCTL). We also showed that the qualitative fragment of the logic can be solved in coNP for the UMC interpretation, and can be solved in polynomial time for a sub-class of UMCs. We also proved lower bounds for these model checking problems. We showed that the model checking problem of IDTMCs with LTL formulas can be solved for both UMC and IMDP semantics by reduction to the model checking problem of IDTMC with omega-PCTL formulas

.Equivalence of labeled Markov chains,
EPFL and CVF

We considered the equivalence problem for labeled Markov chains (LMCs), where each state is labeled with an observation. Two LMCs are equivalent if every finite sequence of observations has the same probability of occurrence in the two LMCs. We showed that equivalence can be decided in polynomial time, using a reduction to the equivalence problem for probabilistic automata, which is known to be solvable in polynomial time. We provided an alternative algorithm to solve the equivalence problem, which is based on a new definition of bisimulation for probabilistic automata. We also extended the technique to decide the equivalence of weighted probabilistic automata. Then, we considered the equivalence problem for labelled Markov decision processes (LMDPs), which asks given two LMDPs whether for every scheduler (i.e., way of resolving the nondeterministic decisions) for each of the processes, there exists a scheduler for the other process such that the resulting LMCs are equivalent. The decidability of this problem remains open. We showed that the schedulers can be restricted to be observation-based, but may require infinite memory.

Stochastic limit-average games are in EXPTIME

We showed that the value of a finite-state two-player zero-sum stochastic game with limit-average payoff can be approximated to within epsilon in time exponential in a polynomial in the size of the game times polynomial in logarithmic in 1/epsilon.

Value iteration,
EPFL:

We surveyed value iteration algorithms on graphs. Such algorithms can be used for determining the existence of certain paths (modelchecking), the existence of certain strategies (game solving), and theprobabilities of certain events (performance analysis). We classified the algorithms according to the value domain (boolean, probabilistic,or quantitative); according to the graph structure (nondeterministic,probabilistic, or multi-player); according to the desired property of paths (Borel level 1, 2, or 3); and according to the alternation depth and convergence rate of fixpoint computations.

-- Brno --

DeVinE Tool

We have further optimized the DiVinE tool and evaluated its performance and scalability on large-scale parallel systems. For preliminary experiments we have used the Distributed ASCI Supercomputer, a wide-area distributed system for Computer Science research in the Netherlands. As a next step we plan to make large scalability tests on the new cluster in Aaalborg. We have extended the DiVinE tool with I/O efficient verification algorithms, which

allow to exploit parallel hard disks. Quantitative LTL mode-checking has been added to the probabilistic version of DiVinE.

-- IRISA --

Verification of Systems with Queues and Stacks:

Many scientific studies analysed the FIFO channel systems, but none offered a fully satisfying solution. We proposed to tackle this problem within the abstract interpretation framework, by defining some abstract lattices adapted to this kind of systems. We considered systems with a finite alphabet of messages, then more complex systems, with an infinite alphabet of messages. This leads us to define and to study a new kind of automata: the lattice automata. Those automata are also useful for the analysis of programs with a call stack.

Quantitative Model-Checking of One-Clock Timed Automata:

We have defined two relaxed semantics (one based on probabilities and the other one based on the topological notion of largeness) for LTL over infinite runs of timed automata which rule out unlikely sequences of events. We proved that these two semantics match in the framework of single-clock timed automata (and only in that framework), and proved that the corresponding relaxed model-checking problems are PSPACE-Complete. Moreover, we proved that the probabilistic non-Zenoness can be decided for single-clock timed automata in NLOGSPACE.

We consider the quantitative model-checking problem for omega-regular properties: we aim at computing the exact probability that a given timed automaton satisfies an omega-regular property. We develop a framework in which we can compute a closed-form expression for this probability; we furthermore give an approximation algorithm, and finally prove that we can decide the threshold problem in that framework.

Probabilistic Büchi Automata:

Probabilistic Büchi automata (PBA) are finite-state acceptors for infinite words where all choices are resolved by fixed distributions and where the accepted language is defined by the requirement that the measure of the accepting runs is positive. The main contribution of this paper is a complementation operator for PBA and a discussion on several algorithmic problems for PBA. All interesting problems, such as checking emptiness or equivalence for PBA or checking whether a finite transition system satisfies a PBA-specification, turn out to be undecidable. An important consequence of these results are several undecidability results for stochastic games with incomplete information, modelled by partially-observable Markov decision processes and omega-regular winning objectives. Furthermore, an alternative semantics for PBA is discussed where it is required that almost all runs for an accepted word are accepting, which turns out to be less powerful, but has a decidable emptiness problem.

Diagnosis and predictability:

We studied the problem of predicting the occurrences of a pattern in a partially-observed discrete-event system. The system is modeled by a labeled transition system. The pattern is a set of event sequences modeled by a finite-state automaton. The occurrences of the pattern are predictable if it is possible to infer about any occurrence of the pattern before the pattern is completely executed by the system. An off-line algorithm to verify the property of predictability is presented. The verification is polynomial in the number of states of the system. An on-line algorithm to track the execution of the pattern during the operation of the system is also presented. This algorithm is based on the use of a diagnoser automaton.

Diagnosis and control synthesis for infomation flow security:

We have been interested in constructing monitors for the detection of confidential information flow in the context of partially observable discrete event systems. We focus on the case where the secret information is given as a regular language. We first characterized the set of observations allowing an attacker to infer the secret behaviors. We considered the general case where the attacker and the administrator have different partial views of the system. Further, based on the diagnosis of discrete event systems, we provide necessary and sufficient conditions under which detection and prediction of secret information flow can be ensured and a construction of a monitor ensuring this task.

Given a finite transition system and a regular predicate, we also addressed the problem of computing a controller enforcing the opacity of the predicate against an attacker (that partially observes the system), supposedly trying to push the system to reveal the predicate. Assuming that the controller can only control a subset of the events it observes (possibly different

from the ones of the attacker), we showed that an optimal control always exists and provide sufficient conditions under which it is regular and effectively computable. These conditions rely on the inclusion relationships between the observable alphabets of the attacker and the controller and the controllable alphabet.

Modular control synthesis:

In this work sufficient conditions for modular (supervisory) control synthesis are presented which equal global control synthesis. In modular control synthesis a supervisory control is synthesized for each module separately and the supervisory control consists of the parallel composition of the modular supervisory controls. The general case of the specification that is indecomposable and not necessarily contained in the plant language, which is often the case in practice, is considered. The usual assumption that all shared events are controllable is relaxed by introducing two new structural conditions relying on the global mutual controllability condition. The novel concept used as a sufficient structural condition is strong global mutual controllability. The main result uses a weaker condition called global mutual controllability together with local consistency of the specification. An example illustrates the approach.

Integration of verification and testing:

A methodology integrating verification and conformance testing for the formal validation of reactive systems has been proposed. A specification of a system - an extended input-output automaton, which may be infinite-state - and a set of safety properties (“nothing bad ever happens”) and possibility properties (“something good may happen”) are assumed. The properties are first tentatively verified on the specification using automatic techniques based on approximated state-space exploration, which are sound, but, as a price to pay for automation, are not complete for the given class of properties. Because of this incompleteness and of state-space explosion, the verification may not succeed in proving or disproving the properties. However, even if verification did not succeed, the testing phase can proceed and provide useful information about the implementation. Test cases are automatically and symbolically generated from the specification and the properties, and are executed on a black-box implementation of the system. The test execution may detect violations of conformance between implementation and specification; in addition, it may detect violation/satisfaction of the properties by the implementation and by the specification. In this sense, testing completes verification.

Discrete controller synthesis for modular reactive systems

We here focused on the exploitation of particularities of modular reactive systems for the application of discrete controller synthesis (DCS). We have proposed a schema of integration of DCS techniques into the modular compilation of an extended synchronous language. In this extended language, the modularity is expressed by nodes, representing components associated with modular synthesis objectives ; we can then obtain, by application of DCS tools on these components, some synchronous controllers controlling parts of programs. In this framework, we implemented a translation schema of a subset of the Lucid Synchrone language into dynamic systems, for further application of Sigali, as DCS tool. Future work will consist in applying decentralized control methods, together with modular distribution of synchronous programs, in order to obtain automatically, from an annotated synchronous program, a distributed controlled system.

This work is part of the post-doc of Gwenael Delaval funded by Artist 2 and a cooperative work between the pop-art EPI (Inria Grenoble) and the VerTeCs EPI (Inria Rennes).

Control of Infinite Symbolic Transitions Systems under Partial Observation

We have proposed algorithms for the synthesis of memoryless controllers through partial observation of infinite state systems modelled by STS. We provide models of safe controllers both for potentially blocking and non blocking controlled systems. To obtain algorithms for this problems, we use abstract interpretation techniques which provide overapproximations of the transitions set to disable. To our knowledge, with the hypotheses taken, the improved version of our algorithm provides a better solution than what was previously proposed in the literature. Our tool SMACS allowed us to make an empirical validation of our methods to show their feasibility and usability.

This work has been done in cooperation with T. Massart and G. Kalyon (Université libre de Bruxelles, Belgium).

STG symbolic test generation tool:

The tool has been improved and a new version can be downloaded from Inria Gforge: https://gforge.inria.fr/plugins/scmsvn/viewcvs.php/?root=bjeannet

-- VERIMAG --

Logics for programs with integer arrays

Programs with integer arrays poses interesting challenges for the existing methods and tools for software verification. In particular, logics for reasoning about infinite state spaces modeling unbounded arrays are required by e.g. predicate abstraction, abstract interpretation or Hoare-style program proofs. Moreover, push-button verification needs decidable logics in which program properties can be expressed.

We have developped two decidable logics in which universally quantified array properties can be expressed. These logics enhance the expressivity of existing logics by allowing arithmetic comparisons between adjacent elements of arrays (such as difference bounds constraints, or octagonal constraints). The decision procedures for the logics are based on translations to classes of counter automata, for which the emptiness problem (existence of a run leading to some final state) is decidable.

Monitoring Real-Time Properties:

Formal verification is a very ambitious activity due to its exhaustiveness and for this reason testing/simulation is still the most commonly used validation techniques. Nevertheless, testing can be made more formal by employing a precise formal specification logic based on temporal logic, against which simulation traces can be checked. This technique called monitoring or runtime verification is gaining popularity and it does not suffer from the state explosion and other difficulties associated with traditional model checking. Motivated by the verification of analog circuits we have developed new algorithms for monitoring timed and hybrid temporal properties, expressed in the logic MITL, against Boolean and analog signals. And, we have developed a tool that has rised interest in several companies (ST, Freescale, Rambus and Mentor Graphics) and applied it to several case studies.

Synthesis from Bounded-Response Properties:

The problem is synthesizing controllers directly from high-level specification is very challenging and proposed theoretical solutions are prohibitively complex. We have suggested a simpler variant of this problem where we restrict ourselves to bounded-response properties which are equivalent to safety properties and hence do not require complex omega automata and their determinization. We express these properties in the real-time logic MITL and then, by transforming them to past properties we can rely on our previous results we construct a deterministic timed automaton to which synthesis algorithms are applied.

Scheduling Policies for Streams of Structured Jobs:

The need to process efficiently streams of tasks that arrive nondeterministically is a crucial problem in embedded system design. Traditional models of real-time systems are not always appropriate for these situations as they often treat periodic and independent tasks. We have defined a new class of scheduling problems where a request generator (a timed language) models the requests which are themselves, partially-ordered sets of tasks (jobs) that may require different types of resources. On these models we prove some fundamental results concerning schedules and scheduling strategies of bounded backlog and latency [DM08].

Formal Verification of Linear Hybrid Automata with PHAVer:

Linear hybrid automata (LHA) are characterized by piecewise constant bounds on the derivatives. They are of interest in formal verification because their dynamics are so simple that basic operators such as discrete and continuous successor states can be computed with exact integer arithmetic, and relatively efficiently over an infinite time horizon. Our tool PHAVer uses exact polyhedral computations to compute the set of reachable states and to verify equivalence and abstraction between automata using assume/guarantee reasoning. Its characteristic is the ability to conservatively overapproximate polyhedra with polyhedra of substantially lower complexity. It is also able to handle more complex dynamics, a generalized form of piecewise affine dynamics, by overapproximation. On-the-fly, adaptive partitioning allows us to target the accuracy of the overapproximation to relevant parts of the state space. Forward/backward refinement, in which the partitioning is iteratively refined while alternating forward and backward reachability, has allowed us to formally verify oscillation of a nonlinear circuit model with three state variables.

Another verification approach for LHA avoids polyhedral computations altogether. It exploits the fact that for LHA reachability along a given path can be formulated as satisfiability of a conjunction of linear constraints, and can therefore be computed very efficiently using linear programming techniques. Various methods of counter example guided abstraction refinement have been proposed in literature to verify safty. We have generically extended these methods to parameter synthesis.

Quantitative verification of embedded software

We have worked on three research directions concerning quantitative verification of embedded software.

The design and implementation of software-intensive embedded product lines requires dealing with a variety of constantly changing application- and system-dependent quantitative non-functional requirements and constraints that need to be verified throughout the development process. Moreover, because product lines are built upon a set of core services which are improved, customized, extended and integrated to come up with differentiated products, there is a need to resort to component-based approaches. However, many embedded applications (e.g., video compression) are most likely specified in a transformational data-oriented style. The componentization of such applications is therefore deferred to the implementation phase, where performance and platform constraints are taken into account. In [YADZB08] we presented a formally-grounded method to carry on this process. The approach consists in integrating (1) the component-based language and execution engine BIP, and (2) the coordination language and code-generation infrastructure FXML/Jahuel. This enables verification of quantitative properties using the associated BIP tool-suite.

-AADL is an aerospace standard for model-driven design of complex real-time embedded systems. Currently, behavioral properties of AADL models can be specified inside the system description using AADL concepts or outside it using external textual languages, and verified using schedulability analysis or (Time Petri Net-based) model-checking tools. Our work [MOYB08] (1) proposes Visual Timed Scenarios (V TS) as a graphical property specification language for AADL, and (2) devises an effective translation from V TS to Time Petri Nets (TPN) which enables model-checking properties expressed in V TS over AADL models using TPN-based tools integrated into AADL-compliant IDEs (e.g., TOPCASED).

-3. In [3] we have developed a technique to compute symbolic polynomial approximations of the amount of dynamic memory required to safely execute a method without running out of memory, for Java-like imperative programs. Given an initial configuration of the stack and the heap, the peak memory consumption is the maximum space occupied by newly created
objects in all states along a run from it. We over-approximate the peak memory consumption using a scoped-memory management where objects are organized in regions associated with the lifetime of methods. We model the problem of computing the maximum memory occupied by any region configuration as a parametric polynomial optimization problem over a polyhedral domain and resort to Bernstein basis to solve it. We apply the developed tool to several benchmarks.

Compositional Verification for Component-based Systems

We have presenedt a compositional method for the verification of component-based systems described in a subset of the BIP language encompassing multi-party interaction without data transfer. The method is based on the use of two kinds of invariants. Component invariants which are over-approximations of components' reachability sets. Interaction invariants which are constraints on the states of components involved in interactions. Interaction invariants are obtained by computing traps of finite-state abstractions of the verified system. The method is applied for deadlock verification in the D-Finder tool. D-Finder is an interactive tool that takes as input BIP programs and applies proof strategies to eliminate potential deadlocks by computing increasingly stronger invariants. The experimental results on non-trivial examples allow either to prove deadlock-freedom or to identify very few deadlock configurations that can be analyzed by using state space exploration.

Incremental Component-Based Construction and Verfication of a Robotic System

Autonomous robots are designed to perform high level tasks on their own , or with very limited external control. They are needed in situations where human control is either infeasible or not cost-effective. Designing and developing software for an autonomous robot is quite a challenging and complex task.

In [BGLN08], [BIS08] we present an evolution of the LAAS Architecture for Autonomous System and its tool GenoM. This evolution is based on the BIP component based design framework which has been successfully used in other domains (e.g. embedded systems). In this study, we show how we seamlessly integrate BIP in the preexisting methodology. We present the componentization of the functional level of a robot, the synthesis of an execution controller as well as validation techniques for checking essential safety properties. This approach has been integrated in the LAAS architecture and we have performed a number of experiment in simulation but also on a real robot (DALA).

Generating Analog-Clock Real-Time Testers Using Action Refinement Techniques:

In a previous work we proposed a method for generating digital-clock tests for real-time systems using action refinement techniques. We have extended this method for generating analog-clock testers. Analog-clock testers are testers which can observe real-time with precision. Our goal from testing is to check the conformance of a given implementation with respect to a given specification (the model). The main benefit of the method is to save memory space needed to build and to store tests. One important contribution of this work is a simplified way for both modelling and testing real-time systems. We first write a (high-level) simplified version of the model of the system, as an input-output transition system (IOTS) and then we refine it into a more detailed (low-level) model as a timed input-output transition system (TIOTS). This same mechanism applies to the test generation procedure.

Automatic Generation of Path Conditions for Concurrent Timed Systems

We concentrate on the automatic generation of test cases for concurrent real-time systems. In order to test a particular behavior of the system, we generate path conditions for (concurrent real-time) execution paths. Instantiating such path conditions allows us to test the desired path. We do not assume finite state systems. Hence our modeled systems may reference unbounded variables in tests and assignments (when we ignore the particular word length in a given machine). Such a precondition characterizes all the states from which we can execute the path. However, there may be other possible executed paths, due to nondeterministic choice, which can be eliminated by adding further synchronization. The path condition calculation can be used in a model checking search, hunting for a path satisfying a given temporal property. It allows us to verify a procedure or collection of procedures in isolation, without providing initial values. Using the weakest precondition calculation, verification is performed symbolically, or ?for all parameters at once?. The temporal property is translated into an automaton and contributes to calculation of the path condition (i.e., it is a condition for executing a path while satisfying the temporal property).

For the real-time case, we need to generalize the calculation of a path condition, taking into account only the essential conditions to follow a particular path in the execution. For example, if the path is abcd, we may constrain only a to precede b, for being on the same process, c to precede d, again, for being on the other process, and b to precede d, for referring to the same variable. We start with a given path (in the flow chart, or interleaved from different flow charts for concurrent processes) merely from a practical consideration; it is very simple to specify an interleaved execution sequence. However, we look at the essential partial order, which is consistent with real-time constraints, rather than at the total order. We cannot assume that transitions must follow each other, unless this order stems from some sequentiality constraints

(such as transitions belonging to the same process or using the same variable) or from timing constraints. Thus, with the above restrictions, acbd is equivalent to abcd and represents the same (partial order) execution. For untimed systems, there is no difference between the condition for partial order execution and the condition for executing any of the sequences (linearizations) consistent with it. Because of commutativity between concurrently executed transitions, we obtain the same path condition either way. However, when taking time constraints into account, the actual time and order between occurrences of transitions does affect the path condition (which now includes time information).

OFFIS

Verification of collision avoidance protocols

We have identified common principles underlying such protocols in what could be called “design patterns” for collision avoidance protocols, building on the key concepts of criticality functions and safety envelopes, and demonstrated the wide applicability with examples from avionics, rail, and automotive applications.

We have provided methods to explore the design space of such protocols, specifically allowing to analyze the interdependencies between system-parameters (e.g. the maximal response times in reaction to external events to ensure collision freedom) and emergent non-functional parameters of components (such as the maximal breaking force, or maximal communication delay for wire-less channels), based on a parametric hybrid logic.

We have developed a verification methodology for such applications using a variant of the ETCS protocol as running example addressing both derived safety and stability requirements.

Verification on non-linear hybrid systems

We have provided abstraction refinement based approaches to the verification of both discrete time and dense time models of non-linear hybrid systems exploiting robustness.

Verification of hybrid systems with large discrete state spaces

We have stretched the limits of maintaining precise abstraction in the verification of hybrid systems with large discrete state spaces by lifting AIG based verification to AIG(T) based model-checking algorithms for both discrete and dense time models of hybrid systems. We have been developing an abstraction refinement based verification technique for open-loop discrete-time models with large discrete state spaces and demonstrated its scalability on industrial size models.

3. Dissemination via Artist2 Events

Events organised and funded by the Artist2 Network of Excellence have been the main "in person" (as opposed to via the website) means by which the NoE has disseminated knowledge. We plan to continue this in the ArtistDesign Network of Excellence.

Both the Artist FP5 project, and then the Artist2 Network of Excellence has organised and funded a huge number of events on Embedded Systems Design:

- Embedded Systems: Industrial Applications '08 *November 12-13, 2008*
- WS on Multicores: Theory and Practice *October 28th, 2008*
- UML&FM'08 *October 27th, 2008*
- WESE'08: WS on Embedded Systems Education *October 23rd, 2008*
- Workshop on Foundations and Applications of Component-based Design (WFCD'2008) *October 19th, 2008*
- RNTS'08 *October 16-17, 2008*
- ACES^{MB} 2008 *September 29th, 2008*
- ARTIST2 Summer School 2008 in Europe *September 8-12, 2008*
- ARTIST2 South-American School for Embedded Systems 2008 *August 25-29, 2008*
- Artist2 Summer School in China 2008 *July 12-18, 2008*
- MoCC 2008 *July 3-4, 2008*
- WCET'08 *July 1st, 2008*
- OSPERT 2008 *July 1st, 2008*
- MPSoc 2008 *June 23-27, 2008*
- Movep'08 *June 23-27, 2008*
- Real-Time Kernels for Microcontrollers: Theory and Practice *June 23-25, 2008*
- COMES 2008 *June 17-18, 2008*
- Mapping of Applications to MPSoCs *June 16-17, 2008*
- ARTIST2 Graduate Course on: Automated Formal Methods for Embedded Systems 2008 *June 16-24, 2008*
- ARTIST2 Graduate Course on Embedded Control Systems *May 26-30, 2008*
- ArtistDesign Workshop on Design for Adaptivity *May 13-14, 2008*
- DataFlow Modeling for Embedded Systems 2008 *May 5th, 2008*
- APRES'08 *April 21st, 2008*
- SLA++P'2008 *April 5th, 2008*
- UML&AADL'2008 *April 2nd, 2008*
- Scopes 2008 *March 13-14, 2008*
- ARTIST2 Timing Analysis activity meeting 2008 *March 13th, 2008*
- ArtistDesign Automotive Systems Day 2008 *March 12th, 2008*

- ATESST Open Workshop *March 3rd, 2008*
- Synchron 2007 *November 26-30, 2007*
- ARTIST2 meeting on Integrated Modular Avionics *November 12-13, 2007*
- WESE'07: WS on Embedded Systems Education *October 4-5, 2007*
- EmSoft'07 *October 1-3, 2007*
- Embedded Systems Week 2007 *September 30th - October 5th 2007*
- Foundations of Component-based Design *September 30th, 2007*
- Between Control and Software (in honor of Paul Caspi) *September 28th, 2007*
- EPSD 2007 *September 10-14, 2007*
- FOSAD 2007 *September 9-15, 2007*
- First European-SouthAmerican School for Embedded Systems *August 21-24, 2007*
- Artist2 / UNU-IIST School in China - 2007 *August 1-10, 2007*
- UML&AADL'2007 *July 14th, 2007*
- FCC 2007 *July 4-5, 2007*
- CAV 2007 *July 3-7, 2007*
- ARTIST WS: Tool Platforms for ES Modelling, Analysis and Validation *July 1-2, 2007*
- ARTIST2 PhD Course on: Automated Formal Methods for Embedded Systems *June 4-12, 2007*
- 2nd Int'l ARTIST Workshop on Control for Embedded Systems *May 31st - June 1st 2007*
- FMGALS'2007 *May 29th, 2007*
- ARTIST2 Graduate Course on Embedded Control Systems *May 7-11, 2007*
- SCOPES 2007 *April 20th, 2007*
- Towards a Systematic Approach to Embedded System Design *April 20th, 2007*
- IRTAW-13 *April 17-19, 2007*
- HSCC'07 *April 3-5, 2007*
- NeRES 2007 *April 2nd, 2007*
- SLA++P 2007 *March 31st, 2007*
- Real-Time Microcontroller Systems: OSEK Standard and experiments on µcontroller devices *March 26-28, 2007*
- ARCS 2007 *March 12-15, 2007*
- ARTIST2 - MOTIVES 2007 *February 19-23, 2007*
- CASTNESS'07 Workshop and School *January 15-17, 2007*
- CASTNESS'07 Workshop and School *January 15-17, 2007*
- ARTIST2 Workshop on Basic Concepts in Mobile Embedded Systems *December 4-5, 2006*
- Synchron 2006 *November 27th - December 1st 2006*

- ARTIST2 Workshop on Timing Analysis in the Industrial Development Process (ISoLA 2006) *November 17th, 2006*
- MoCC - Models of Computation and Communication *November 16-17, 2006*
- Artist2 - Foundations and Applications of Component-based Design *October 26th, 2006*
- WESE'06 - Embedded Systems Education *October 26th, 2006*
- ATVA China 2006 *October 23-26, 2006*
- ATVA China 2006 *October 23-26, 2006*
- JTRES 2006 *October 11-13, 2006*
- MARTES 2006 *October 2nd, 2006*
- ADSD 2006: Advanced Digital Systems Design *September 25-29, 2006*
- FOSAD 2006: 6th International School on Foundations of Security Analysis and Design *September 10-16, 2006*
- First European Laboratory on Real-Time and Control for Embedded Systems *July 10-14, 2006*
- CORDIE'06: Concurrency, Real-Time and Distribution in Eiffel–like Languages *July 4-5, 2006*
- ARTIST2 Workshop on Requirements for Flexible Scheduling in Complex Embedded Systems *June 16th, 2006*
- ARTIST2 Workshop on Execution Platforms / Cluster Meeting *May 22-23, 2006*
- ARTIST2 Workshop on Specification and Verification of Secure Embedded Systems *May 18th, 2006*
- ARTIST2 / UNU-IIST Spring School in China 2006 *April 3-15, 2006*
- ARTIST2 Graduate Course on Embedded Control Systems *April 3-7, 2006*
- ARTIST2 Workshop Beyond AutoSar *March 23-24, 2006*
- ARTIST Workshop at DATE'06 *March 10th, 2006*
- Workshop: Distributed Embedded Systems *November 21-24, 2005*
- ARTIST2 Summer School 2005 *September 29th - October 2nd 2005*
- WESE'05 - ARTIST2 Workshop on Embedded Systems Education *September 22nd, 2005*
- 31st EUROMICRO Conference - Special session: Model Driven Engineering (MDE) *August 30th - September 3rd 2005*
- ACM-IEEE MEMOCODE'2005 *July 11-14, 2005*
- IST/NSF: Transatlantic Research Agenda on Future Challenges in Embedded Systems Design *July 8th, 2005*
- EU/US: Component-based Engineering for Embedded Systems *July 7th, 2005*
- OSPERT 2005 *July 5th, 2005*
- ARTIST Seminar on Adaptive Real-Time Systems *June 20-23, 2005*
- ARTIST Workshop at DATE'05 *March 11th, 2005*
- HSCC '05 - Hybrid Systems: Computation and Control *March 9-11, 2005*
- First S.Ha.R.K. Workshop *February 28th - March 4th 2005*

4. Dissemination via the Artist2 Web Portal

The Artist2 Web Portal (http://www.artist-embedded.org/) is the principal means by which information is dissemainted by the Artist2 NoE within the Embedded Systems Community.

The Artist2 Web Portal will continue to evolve and be maintained as the ArtistDesign Web Portal (2008 – 2010).

4.1 Objectives and Background Information

The aim of the Artist2 Web Portal is 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 (Artist2 core and affiliated partners), the Artist2 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 Artist2 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) is maintained automatically.
- Ergnomics are set for the entire site. The "look and feel" of the site is always homogeneous thoughout the site. It's possible to change these ergonomics, and these changes are applied homogeneously throughout the site, via automated machanisms.

4.2 Google keywords used to access the site

A representative sample of recent google searches, used to access the site, include the following. It's very interesting to note how well the portal is placed, when searches on topics not directly related to the NoE are used.

NB1: These links are active, so you can see the results of the google search yourself by clicking on the keywords.
NB2: Over time, google results will shift.

- « castness »
- « hipeac »
- « paul caspi retirement »
- « "semantic level" component »
- « artist2 »
- « cluster testing »
- « FPVI project » (2)
- « EMBEDDED C »
- « spreading excellence in research »
- « presentation material »
- « ERTS 2009 embedded »
- « project conclusion application design presentation »
- « ERTS EMBEDDED REAL TIME SOFTWARE 2009 »
- « Towards a traceability model in a MARTE-based methodology for real-time embedded systems »
- « Edward A Lee artist »
- « what is cluster testing » (2)
- « main projects » (2)
- « as level art brief example »
- « neeraj suri »
- « modular avionics design »
- « timing analysis embedded »
- « integrated modular avionics »
- « established trust level EAL7 »
- « Industrial Phd Aalborg »

- « cédric di tofano »
- « ARTIST2 Design »
- « Multi-Clock Latency-Insensitive Architecture »
- « Real Time Components »
- « Systems integration activities »
- « Integrated Modular Avionics »
- « ahmed bouajjani »
- « rtcsa 2009 »
- « FORMAL ANALYSIS ART TUTORIAL »
- « NOE structure »
- « RTSS 2009 »
- « artistdesign »
- « Embedded Systems Week 2009 »
- « UML QoS profile »
- « Suitability of dynamic load balancing in resource constrained embedded systems: An overview of challenges and limitations »
- « esweek 2009 »
- « embedded software automotive memory slides »
- « integrated modular avionics »
- « t »
- « survey of programming languages »
- « survey of programming language »
- « EMSOFT 2009 »
- « cluster testing »
- « On the Scalability of Real-Time Scheduling Algorithms on Multicore Platforms: A Case Study »
- « artist design artist2 »
- « artist »
- « "abhijit datar" »
- « artist embedded »
- « artist2 »
- « integrated modular avionics »
- « Automation System links.html »
- « erts »
- « ROBERT BOSCH AG »
- « Integrated modular avionics for first time »
- « Dr. Christian Salzmann BMW »
- « parallel optimization demo platform compiler options »
- « first to use Integrated Modular Avionics (IMA) »
- « Sirena "Real time Embedded Networked Applications" »
- « integrated modular avionic »
- /custom?hl=en&client=pub-899...
- « hardware course material online »
- « nanjing university yue zhang »
- « processor-memory co-exploration driven by memory-aware architecture »
- « mpsoc 2009 »
- « "PARADES" Roma company »
- « real-time embedded system »
- « university of uppsala embedded systems »
- « telecommunication workshops seminars in 2009 »
- « organizational chart of distributed embedded system »
- « Real Time Application Support (AEP), IEEE Std 1003.13-1998 »
- « bound-t »
- « Ptolemy Project »
- « open PhD Positions in Denmark »
- « artistdesign »
- « phan thi xuan linh »
- « artist »
- « artist embedded »
- « artist of this time »
- « artist noe »
- « AADL VTS »
- « quantitative testing »
- « ARTIST DESIGN embedded »
- « 2009 embedded system »
- « universities in america offering embedded systems course »
- « artemis fp7 »
- « WCPS2008 »
- « summer Engineering courses in Switzerland »
- « Languages and Tools for Hybrid Systems Design »
- « verification assembler code »
- « basics of model checking paul gastin »
- « combest passerone »
- « "roberto zafalon" »
- « security protocol term rewriting »
- « ISO/IEC TR 18037:2004 »
- « Device Power Scheduling via Systems Management »
- « Ndukwu northeastern »
- « artist2 fpga »
- « verification testing tools »
- « CODES-ISSS 2009 »
- « adaptive real-time »
- « truetime lecture »
- « truetime matlab »
- « school in europe »
- « osek oil »
- « what is clusters in strategic management »
- « powerlink matlab toolbox »
- « "eric lenormand" -vin -vins -théatre -pièce -costumes -décors »
- « hardware VLSI research professor .edu home page »
- « international collaboration »
- « uml real time »
- « computing design »
- « Prof. Francky Catthoor »
- « Features of Real-Time Embedded Systems »
- « mailing list artists »

4.3 Structure of the Web Portal

The structure of the Artist2 web site at the end of Year 3 is as follows (visible on the Site Map: http://www.artist-embedded.org/artist/spip.php?page=plan).

The links below are active.

About the Artist2 NoE

- Strategic Objectives
- Approach
- Joint Programme of Activities (JPA)
- Artist2 Core Partners
- Workshops
- Workshops and Conferences
- Education
  - Educational Methods for Embedded Systems Design
  - Events and Publications on Specific Topics
- International Collaboration
- Contributions to Standards
- State of the Art
- Related Projects
- Becoming an Affiliated Partner
- Site Map

Participants

- **Artist2 Participants**
  - Strategic Management Board — Artist2 NoE
  - Artist2 Core Partners - full consortium
  - Artist2 Core Partners - by Cluster topics
    - Cluster: Real-Time Components
    - Cluster: Adaptive Real-Time
    - Cluster: Compilers and Timing Analysis
    - Cluster: Execution Platforms
    - Cluster: Control for Embedded Systems
    - Cluster: Testing and Verification
  - Artist2 Affiliated Partners
    - Affiliated Industrial Partners
    - Affiliated SME Partners
    - Affiliated Academic Partners
    - Affiliated International Collaboration Partners
- **Core Team Leaders**

Research and Integration

- **Cluster: Real-Time Components**
  - Research and Integration Activities for the "Real Time Components" cluster
- **Cluster: Adaptive Real-Time**
  - Research and Integration Activities for the "Adaptive Real Time" cluster
- **Cluster: Compilers and Timing Analysis**

- Research and Integration Activities for the "Compilers and Timing Analysis" cluster
- **Cluster: Execution Platforms**
  - Research and Integration Activities for the "Excution Platforms" cluster
  - Internal Meetings
    - Execution Platforms Cluster Meeting
- **Cluster: Control for Embedded Systems**
  - Research and Integration Activities for the "Control for Embedded Systems" cluster
- **Cluster: Testing and Verification**
  - Research and Integration Activities for the "Testing and Verification" cluster

Dissemination

- **Workshops**
  - ATESST Open Workshop
  - MARTES 2008
  - ArtistDesign
  - SLA++P'2008
  - SIES'2008
  - UML&AADL'2008
  - Scopes 2008
  - APRES'08
  - WCPS2008
  - WTR 2008
  - Workshop on Foundations and Applications of Component-based Design (WFCD'2008)
  - Mapping of Applications to MPSoCs
  - WCET'08
  - MPSoc 2008
  - ArtistDesign Workshop on Design for Adaptivity
  - DataFlow Modeling for Embedded Systems 2008
  - ARTIST2 Timing Analysis activity meeting 2008
  - MoCC 2008
  - ACES^{MB} 2008
  - UML&FM'08
  - WESE'08: WS on Embedded Systems Education
  - Movep'08
  - IMCSIT'08 - Real Time Systems Workshop
  - OSPERT 2008
  - Embedded Control Systems: From Design to Implementation
  - COMES 2008
  - WS on Multicores: Theory and Practice
  - Mapping Applications to MPSoCs 2009
  - Embedded Systems: Industrial Applications '08
  - SCOPES 2009
  - Optimizations for DSP and Embedded Systems 2009
  - DySCAS 2009
  - Workshops and Seminars in 2007
    - COCV 2007
    - SEUS 2007
    - UML&AADL'2007

- SCOPES 2007
- CASTNESS'07 Workshop and School
- DCDS'07
- SIES'2007
- IRTAW-13
- Towards a Systematic Approach to Embedded System Design
- Distributed Object Computing for RT and Embedded Systems
- NeRES 2007
- Software Tools for Multi-Core Systems
- ARTIST2 meeting on Integrated Modular Avionics
- SLA++P 2007
- WPDRTS 2007
- FMGALS'2007
- LCTES'07
- Dagstuhl: Geometry in Sensor Networks
- Dagstuhl: Mobile Interfaces Meet Cognitive Technologies
- Dagstuhl: Tools for the Model-based Development of Certifiable, Dependable Systems
- Dagstuhl: Model-Based Engineering of Embedded Real-Time Systems
- Dagstuhl: Formal Protocol Verification Applied
- FCC 2007
- ARTIST WS: Tool Platforms for ES Modelling, Analysis and Validation
- WCET'07
- Between Control and Software (in honor of Paul Caspi)
- Synchron 2007
- Precise Behavioral Semantics for DSML
- WESE'07: WS on Embedded Systems Education
- Foundations of Component-based Design
- 2nd Int'l ARTIST Workshop on Control for Embedded Systems

- Workshops and Seminars in 2006
  - CORDIE'06: Concurrency, Real-Time and Distribution in Eiffel–like Languages
  - Artist2 - Foundations and Applications of Component-based Design
  - MARTES 2006
  - JTRES 2006
  - WESE'06 - Embedded Systems Education
  - ARTIST2 Workshop on Timing Analysis in the Industrial Development Process (ISoLA 2006)
  - MoCC - Models of Computation and Communication
  - ARTIST2 Workshop on Requirements for Flexible Scheduling in Complex Embedded Systems
  - ARTIST2 Workshop on Specification and Verification of Secure Embedded Systems
  - ARTIST2 Workshop Beyond AutoSar
  - ARTIST Workshop at DATE'06
  - ARTIST2 Workshop on Execution Platforms / Cluster Meeting
  - ARTIST2 Workshop on Basic Concepts in Mobile Embedded Systems
  - Synchron 2006
  - ATVA China 2006
  - ATVA China 2006
- Workshops and Seminars in 2005
  - ACM-IEEE MEMOCODE'2005
  - Workshop: Distributed Embedded Systems

    - WESE'05 - ARTIST2 Workshop on Embedded Systems Education
    - OSPERT 2005
    - ARTIST Seminar on Adaptive Real-Time Systems
    - ARTIST Workshop at DATE'05
    - HSCC '05 - Hybrid Systems: Computation and Control
    - First S.Ha.R.K. Workshop
    - EU/US: Component-based Engineering for Embedded Systems
    - IST/NSF: Transatlantic Research Agenda on Future Challenges in Embedded Systems Design
    - 31st EUROMICRO Conference - Special session: Model Driven Engineering (MDE)

- **Schools and Seminars**
    - ADSD 2006: Advanced Digital Systems Design
    - First European Laboratory on Real-Time and Control for Embedded Systems
    - First European-SouthAmerican School for Embedded Systems
        - First European-SouthAmerican School for Embedded Systems - Programme
    - ARTIST2 Graduate Course on Embedded Control Systems
    - FOSAD 2006: 6th International School on Foundations of Security Analysis and Design
    - ARTIST2 - MOTIVES 2007
        - Social Event
    - ARTIST2 / UNU-IIST Spring School in China 2006
    - ARTIST2 Graduate Course on Embedded Control Systems
    - ARTIST2 Summer School 2005
    - Artist2 / UNU-IIST School in China - 2007
    - MDD4DRES
    - ARTIST2 South-American School for Embedded Systems 2008
        - School posters
    - CASTNESS'07 Workshop and School
    - Quantitative Aspects of Embedded Systems
    - FOSAD 2007
    - ARTIST2 Graduate Course on Embedded Control Systems
    - Real-Time Microcontroller Systems: OSEK Standard and experiments on µcontroller devices
    - EPSD 2007
    - ARTIST2 PhD Course on: Automated Formal Methods for Embedded Systems
    - LASER Summer School on Software Engineering
    - CASTNESS 2008
    - ARTIST2 Summer School 2008 in Europe
    - Artist2 Summer School in China 2008
    - ARTIST2 Graduate Course on: Automated Formal Methods for Embedded Systems 2008
    - Real-Time Kernels for Microcontrollers: Theory and Practice
    - MDD for Distributed Real-time Embedded Systems (MDD4DRES) 2009
    - FOSAD 2009
- **International Collaboration**
- **Publications**
- **Contributions to Standards**
    - Modelling

    - Programming Languages
        - ARTIST Survey of Programming Languages
    - Operating Systems and Middleware
- **Course Materials Available Online**

Embedded System Links

- **Journals**
- **Conferences**
    - MEMOCODE 2007
    - EmSoft'07
    - DAC 2007
    - DATE 2007
    - RTAS 2008
    - CODES+ISSS 2006
    - IST Event 2006
    - RTSS 2006
    - FM 2006
    - CASES 2007
    - ASP-DAC 2008
    - HSCC'07
    - ARCS 2007
    - ECRTS 2007
    - IESS'07
    - ECMDA
    - ESEC/FSE
    - ECC
    - FDL'07
    - CAV 2007
    - SAMOS VII
    - RTSS 2007
    - ETFA 2007
    - RTS 2007
    - Networks-on-Chip Symposium
    - RTNS'2007
    - FORMATS'07
    - Embedded Systems Week 2007
    - Embedded Systems Conference 2007
    - RTCSA 2007
    - CODES-ISSS 2007
    - ECRTS 2008
    - LCTES'08
    - DATE'08
    - ERTS 2008
    - Ada-Europe'08
    - IFAC'08
    - RNTS'08
    - ESWEEK 2008
    - Cyber Physical Systems Week 2008
    - Models'08
    - DCOSS '08

    - CAV 2009
    - DATE 2009
    - Cyber Physical Systems Week 2009
    - LCTES'09
    - POPL 2009
    - ESWeek 2009

- **Hot Topics**
- **Standards**
- **Tools and Platforms**
    - Real-Time Components
    - Adaptive Real-Time
    - Compilers and Timing Analysis
    - Control for Embedded Systems
    - Testing and Verification
- **Main Projects**
    - ARTEMIS European Technology Platform
- **Position Papers**
- **Roadmaps**
- **Newsletters and Magazines**
- **Mainstream Press**
- **Announcements**
    - Artist Mailing List
    - Open Positions in Embedded Systems
        - Previous Open Positions
    - Other Calls
    - Other
- **Publications**

5. Publications

All of the following types of publications are available via the Artist Web Portal.

http://www.artist-embedded.org/

5.1 Major Surveys, Textbooks and Roadmaps

We cite here some of the major studies, surveys, and roadmaps from both Artist FP5 (the first Artist project) and the Artist2 NoE.

5.1.1 Artist2 Survey of Programming Languages

http://www.artist-embedded.org/artist/ARTIST-Survey-of-Programming.html

Alan Burns (Editor)

The production of real-time and embedded systems involves the use of many different tools and techniques. As these systems become more software centric, programming languages employed in the production of this software are now of crucial importance. Any language has the dual role of enabling expression whilst at the same time limiting the framework of concept and abstractions within which that expressive power may be applied. If a language does not support a particular notion then programmers cannot apply it and may even be totally unaware of its existence. For real-time programming languages there are many such concepts that are supported to a greater or lesser extent in a range of languages that purports to be appropriate for the embedded systems domain. For example: time, clocks, concurrency, deadlines, events and signals, exceptions, periodicity, scheduling and predictability are all important notions that programmers may wish to address and which should therefore be available to them via the implementation languages that they can employ.

This survey considers over twenty programming languages. The short reports available on each language aim to introduce the main features of the language, provide the links to further sources of information, and give an indication of the current developments within the language. Some languages are mature, used widely and are the subject of rigorous standardisation procedures. Others are research languages, with a small user population and an informal definition. All languages covered in the survey are implementation languages in that they are supported by tools (typically compilers) which generate executable code for the designated hardware platform. To give a structure to the survey each language is placed in one of five classes: imperative, functional, synchronous, model-based, and platform-based. However this is a loose classification as some languages could easily be placed in more than one category, and imperative languages can usually be constrained to mimic the other styles.

A final point to note about this survey is that it is inevitably incomplete and it is a dynamic document. Extra contributions can easily be added (email the editor). Similarly, additional surveys of the languages that are covered can be included.

The report covers :

- Imperative Languages
- Functional Languages
- Synchronous Languages
- Model-Based Languages
- Platform-Based Language

5.1.2 Artist2: Languages and Tools for Hybrid Systems Design

http://www.artist-embedded.org/artist/Languages-and-Tools-for-Hybrid,442.html

The explosive growth of embedded electronics is bringing information and control systems of increasing complexity to every aspect of our lives. The most challenging designs are safety-critical systems, such as transportation systems (e.g., airplanes, cars, and trains), industrial plants and health care monitoring. The difficulties reside in accommodating constraints both on functionality and implementation. The correct behaviour must be guaranteed under diverse states of the environment and potential failures; implementation has to meet cost, size, and power consumption requirements. The design is therefore subject to extensive mathematical analysis and simulation. However, traditional models of information systems do not interface well to the continuous evolving nature of the environment in which these devices operate. Thus, in practice, different mathematical representations have to be mixed to analyze the overall behaviour of the system. Hybrid systems are a particular class of mixed models that focus on the combination of discrete and continuous subsystems. There is a wealth of tools and languages that have been proposed over the years to handle hybrid systems. However, each tool makes different assumptions on the environment, resulting in somewhat different notions of hybrid system. This makes it difficult to share information among tools. Thus, the community cannot maximally leverage the substantial amount of work that has been directed to this important topic. In this paper, we review and compare hybrid system tools by highlighting their differences in terms of their underlying semantics, expressive power and mathematical mechanisms. We conclude our review with a comparative summary, which suggests the need for a unifying approach to hybrid systems design. As a step in this direction, we make the case for a semantic-aware interchange format, which would enable the use of joint techniques, make a formal comparison between different approaches possible, and facilitate exporting and importing design representations.

5.1.3 Artist2: Tools for Real--Time Control Systems Codesign

http://www.artist-embedded.org/artist/Tools-for-Real-Time-Control,445.html

Authors: Dan Henriksson, Ola Redell, Jad ElKhoury, Martin Törngren, and KarlErik Årzén: Department of Automatic Control Lund Institute of Technology April 2005

This report presents a survey of current simulation tools in the area of integrated control and real- time systems design. Each tool is presented with a quick overview followed by a more detailed section describing comparative aspects of the tool. These aspects describe the context and purpose of the tool (scenarios, development stages, activities, and qualities constraints being addressedE and the actual tool technology Dtool architecture, inputs, outputs, modeling content, extensibility and availability).

The tools presented in the survey are the following; Jitterbug and TrueTime from the Department of Automatic Control at Lund University, Sweden, AIDA and XILO from the Department of Machine Design at the Royal Institute of Technology, Sweden, Ptolemy II from the Department of Electrical Engineering and Computer Sciences at Berkeley, California, RTSIM from the RETIS Laboratory, Pisa, Italy, and Syndex and Orccad from INRIA, France.

The survey also briefly describes some existing commercial tools related to the area of real-time control systems.

5.1.4 Artist2 Roadmap on Control of Real-Time Computing Systems

Written by the cluster: Control for Embedded Systems Cluster, 2005
http://www.artist-embedded.org/docs/Events/2005/Artist2_Y1Review/Deliverables/Artist2/18b_Control_Roadmap.pdf

An important result of the EU ARTIST FP5 project was four roadmaps on Hard Real-Time Development Environments, Component-Based Design and Implementation Platforms, Adaptive Real-Time Systems for Quality of Service Management, and Execution Platforms respectively, [Bouyssounouse and Sifakis, 2005]. The current roadmap written by the partners of the Control for Embedded Systems cluster within the EU/IST FP6 Network of Excellence ARTIST2 can partly be viewed as an extension of the adaptive real-time system roadmap. The focus is how flexibility, adaptivity, performance and robustness can be achieved in a real-time computing or communication system through the use of control theory.

Similar to the ARTIST roadmaps this roadmap is intended as a roadmap for research rather than an roadmap on industrial R & D in general. The roadmap is **not** a roadmap on real-time control. In real-time control the real-time computing system is used as an implementation platform for a control system controlling some external dynamical system, often a physical plant with external inputs and outputs. Here, it is instead the real-time computing system that is the subject to the control. The item that is controlled is in most cases the allocation of computing and communication resources, e.g., the distribution or scheduling of CPU time among different competing tasks, jobs, requests, or transactions. Due to this, control of computing systems also goes under the name of feedback scheduling.

The roadmap assumes basic knowledge in real-time computing and control engineering from the readers. In parallel to the current roadmap a separate roadmap on Real-Time Techniques in Control System Implementation has been developed.

5.1.5 "Embedded System Design" textbook by Peter Marwedel, TU Dortmund

http://www.artist-embedded.org/artist/Embedded-Systems-Design,441.html

Embedded systems can be defined as information processing systems embedded into enclosing products such as cars, telecommunication or fabrication equipment. Such systems come with a large number of common characteristics, including real-time constraints, and dependability as well as efficiency requirements. Following the success of information technology (IT) for office and workflow applications, embedded systems are considered to be the most important application area of IT during the coming years. This importance of embedded systems is so far not well reflected in many of the current curricula. Embedded System Design is intended as an aid for changing this situation. It provides the material for a first course on embedded systems, but can also be used by PhD students and professors.

A key goal of this book is to provide an overview of embedded system design and to relate the most important topics in embedded system design to each other. It should help to motivate students as well as professors to put more emphasis on education in embedded systems.

5.1.6 ARTIST FP5 Roadmap for Embedded Software and Systems

http://www.artist-embedded.org/artist/ARTIST-FP5-Roadmap-for-Embedded.html

This extensive and increasing use of embedded systems and their integration in everyday products mark a significant evolution in information science and technology. Nowadays embedded systems design is subject to seamless integration with the physical and electronic environment while meeting requirements like reliability, availability, robustness, power consumption, cost, and deadlines. Thus, embedded systems design raises challenging problems for research, such as security, reliable and mobile services, large-scale heterogeneous distributed systems, adaptation, component-based development, and validation and tool-based certification.

This book results from the ARTIST FP5 project funded by the European Commission. By integration 28 leading European research institutions with many top researchers in the area, this book assesses and strategically advances the state of the art in embedded systems. The

coherently written monograph-like book is a valuable source of reference for researchers active in the field and serves well as an introduction to scientists and professionals interested in learning about embedded systems design.

5.1.7 Artist FP5 / ACM Transactions in Embedded Computing Systems Special Issue on Education

http://www.artist-embedded.org/artist/ACM-Special-Issue-on-Education.html

Embedded systems applications now include a very large proportion of the advanced products designed in the world, spanning transport (avionics, space, automotive, trains), electrical and electronic appliances (cameras, toys, television, washers, dryers, audio systems, and cellular phones), process control (energy production and distribution, factory automation and optimization), telecommunications (satellites, mobile phones and telecom networks), and security (e-commerce, smart cards), etc.. The relative weight of software in the value of embedded systems is constantly expanding. The extensive and increasing use of embedded systems and their integration in everyday products marks a significant evolution in information science and technology.

There is now a strategic shift in emphasis for embedded systems designers: from simply achieving feasibility, to achieving optimality. Optimal design of embedded systems means targeting a given market segment at the lowest cost and delivery time possible. Optimality means seamless integration with the physical and electronic environment while respecting real-world constraints such as hard deadlines, reliability, availability, robustness, power consumption, and cost. In our view, optimality can only be achieved through the emergence of embedded systems as a discipline in its own right.

An important factor for the emergence of embedded systems as a discipline is the existence of integrated curricula for training engineers and researchers, able to tackle a range of topics which until now had been spread across many different areas, including: general computer science and engineering, real-time computing, systems architecture, control and signal processing, security and privacy, networking, mathematics, electronics.

This special issue of the ACM Transactions in Embedded Computing Systems aims to provide the basis for integrated undergraduate and graduate curricula covering the essential areas of knowledge for tomorrow’s embedded systems engineers and researchers.

5.1.8 Artist FP5 Guidelines for a Graduate Curriculum on Embedded

http://www.artist-embedded.org//docs/Publications/Education.pdf

The design of real-time embedded systems requires skills from three specific disciplines: control theory, computer science, and electronic engineering, and their combination. This often involves experts from differing backgrounds, who do not recognize that they address different issues from complementary angles.

The motivation for defining a specific curriculum in embedded systems is mainly to promote the training of engineers with expertise in the above three disciplines.

5.2 Artist Mailing List

The Artist Mailing List is used for widespread dissemination of short announcements about events, calls for papers, open positions and other announcements in Embedded Systems Design. Several announements are disseminated on the Artist Mailing List each week.

Anyone may subscribe to the Artist Mailing List.

Anyone who has subscribed to the Artist Mailing List may post announcements to it, although certain restrictions (see belox) apply. The list is moderated by the Artist Technical Coordinator.

The Artist Mailing List is separate from the one for the Artist Newsletter.

How to Subscribe

To subscribe or unsubscribe to the Artist Mailing List, the subscriber simply sends an appropriate message to the mail server, then *confirms his/her identity* by replying to the message sent by the mail server.

New subscribers receive a welcome message confirming their subscription to the Artist Mailing List.

Submitting an Announcement

To submit an announcement, a subscriber simply sends their message to: announcements@lists.artist-embedded.org.

Once the moderator has approved it, the message is sent to the mailing list as is.

Restrictions on Messages

To ensure that interest in the list remains high, the list is moderated to check that these few common-sense guidelines are followed:

1. Only announcements of **interest to the general community**, of a non-commercial nature, will be accepted.
2. Please plan to submit **only one announcement** for a given event.
3. Attachments should be **very small**. If you want to provide larger information, make it available on a website and just provide the link.
4. The **name of the person** sending the message should appear clearly in the "From" field (eg: Fred Astaire <fastaire@domain.uk>).
5. The topic must appear clearly on the subject line (preferably just the title of the event or announcement)
6. The **Body of the message** must contain the main elements. Details may be provided in an attachment, but it's not OK to have only an attachment.
7. A **Subject line** is needed. It's best to put only the full title of the event or announcement.

Privacy Policy

Artist has a very strict privacy policy: we will not give the list of subscribers to any outside parties, including the Artist core participants.

5.3 Videos

In Year4, the Artist2 Summer school has started recording videos of the lectures, which will be made available free of charge via the Artist web portal.

At the time of this writing, the videos were still in post-production, but a number of these are ready here:

- **Klaus Havelund** (NASA JPL)
  Rule-based Runtime Verification
- **Giorgio Buttazzo** (Pisa)
  Real-Time Scheduling and Resource Management
- **Reinhard Wilhelm** (Saarland)
  Timing Analysis and Timing Predictability
- **Kai Richter** (Symtavision)
  Establishing Formal Scheduling Analysis in Automotive Design Processes
- **Karl Erik Arzen** (Lund)
  Pedro Albertos (UP Valencia)
  Implementation of control systems in resource-constrained embedded systems
- **Ed Brinksm** (ESI)
  Quantitative Testing Theory
- **Michael Gonzalez** (U. Cantabria)
  Contract-based resource reservation and scheduling
- **Lothar Thiele** (ETHZ)
  Performance analysis of distributed real-time systems
- **Diederik Verkest** (IMEC)
  Mapping C code on MPSoC for Nomadic Embedded Systems
- **Luis Almeida** (U. Aveiro)
  The challenges of flexible real-time communication
- **Kim Larsen** (Aalborg)
  Quantitative Verification and Synthesis for Embedded Systems
- **Hermann Haertig** (TU Dresden)
  Enforceable Component-Based Realtime Contracts
- **Peter Marwedel & Heiko Falk** (TU Dortmund)
  Memory architecture aware compilation
- **Rance Cleaveland** (University of Maryland, USA)
  An instrumentation -based Approach to Controller Model Validation
- **Raj Rajkumar** (Carnegie Mellon University)
  Building Blocks for Large-Scale Wireless Sensor Networks

- **Marta Kwiatkowska** (Oxford U.)
  Quantitative Probabilistic Verification of Systems

- **Gerard Berry** (Esterel Technologies)
  The evolution of the synchronous model

- **Marco Bekooij** (NXP)
  Dataflow analysis for predictable multiprocessor design

- **Steve Vestal** (Honeywell)
  Automating compositional safety analysis for IMA systems

An overview video clip of the school was also made:

- **ARTIST2 Summer School 2008 in Europe**
  http://artist.imavox.ch/clip_autrans/

5.4 Research Publications

The Artist2 community has been very active in publishing in scientific journals and conferences, as attested by the list of publications provided in this document. Clearly, this represents a huge amount of work. Publication of research is a bottom-up process, which may seem chaotic – but this is intrinsic to research.

The lists of references for joint papers alone (research papers authored by representatives from two or more Artist2 partners), across the 4 years of the NoE, span 43 pages.

5.5 Newsletter

The Artist2 Network of Excellence has also published a high-quality Newsletter, which will be continued into ArtistDesign. http://www.artist-embedded.org/artist/Artist2-Newsletter,438.html

The following pages show a sample edition of the Newsletter.

The newletter is distributed electronically to a mailing list with over 20000 entries.

ARTIST Newsletter n°4

ARTIST Newsletter n°4
September 2007

Newsletter Subscription

Subscription to this newsletter is free of charge. If you would like to be added or removed from the subscription list, simply send a message to the editor:

Bruno.Bouyssounouse@imag.fr

Online

This newsletter is also available online:
http://www.artist-embedded.org/artist/Artist2-Newsletter-438.html

Overview of this issue

This issue includes:

- descriptions of the main upcoming European workshops, conferences and schools on embedded systems design;
- invited articles from EADS on «Component-based Design and Implementation of Avionics Systems» and from Nokia on «Performance / Power Measurements for Embedded Multi-Core Processors»;
- a series of focus articles on Adaptive Real Time.

About Artist2

The long-term objective of the ARTIST2 Network of Excellence on Embedded Systems Design (http://www.artist-embedded.org/) is to build a durable European research community on Embedded Systems Design, by integrating the topics, teams, competencies, from 7 essential topics: Real-Time Components, Adaptive Real-Time, Compilers and Timing Analysis, Execution Platforms, Control for Embedded Systems, Testing and Verification.

The overall objective of Artist2 is the emergence of Embedded Systems Design as a mature scientific and engineering discipline, through tight integration of central players from the European research community. A central mission for the NoE is to disseminate excellence in the area, through an ambitious Joint Programme of Activities for Spreading Excellence, including Education and Training, Dissemination and Communication, Industrial Liaison, and International Collaboration.

Artist2 has 40 partners, including 2 large industrial and 3 SMEs. It integrates joint research activities at two levels:

- *Integration within clusters*. Currently, efforts in the area are fragmented, and no single European research group gathers sufficient critical mass. Integrating the clusters is the first step towards integrating the area as a whole.
- *Integration between cluster topics* to create the multi-disciplinary community that will pilot the embedded systems design area. This will be achieved through research and integration activities that will bring together teams from different clusters.

Expected Results

ARTIST2 has a durable structuring effect on European research:

- There is a direct impact on the integration of academic research. It allows for new, coherent theoretical frameworks to emerge, particularly those that can contribute to the unification of the area. For this, the NoE takes measures to overcome the inherent contextual, cultural, and disciplinary diversity.
- ARTIST2 will impact R&D activity from an organizational perspective. ARTIST2 explicitly aims to create a context, an infrastructure and a culture for the design of joint, multi-organisational, multi-disciplinary R&D work in embedded systems design.
- ARTIST2 has a structural impact on European education in Embedded Systems Design, by:
  - Integrating state of the art knowledge into the curricula and accelerating convergence towards multi-disciplinary approaches.
  - Promoting approaches and techniques, which are well-adapted to meeting current and future industrial needs.

Table of Contents

Mailing List

To receive *frequent* information on topics in Embedded Systems Design (workshops, open positions, etc), subscribe to our mailing list:

http://www.artist-embedded.org/artist/Artist-Mailing-List.html

This list is separate from the one for the Artist newsletter.

Workshops

Newsletter
September 2007

Artist Web Portal

The list of upcoming ARTIST-related events is available online:
http://www.artist-embedded.org/

MARTES'07

Modeling and Analysis of Real-Time and Embedded Systems
Oct 2nd 2007
Nashville, Tenessee (USA)
with MoDELS/UML 2007
organised with Artist partners

This workshop addresses all aspects of the representation, analysis, and implementation of DRES models, related to, but not limited to, the following principal topics:

- Modeling RT/E using modeling languages such as UML
- Semantic aspects of real-time in UML and similar modeling languages
- Methods and tools for analysis of RT systems and components

http://www.artist-embedded.org/artist/-MARTES-2007-.htmlWCET-07.html

Between Control and Software

In honor of Paul Caspi
September 28th, 2007
VERIMAG - Grenoble, France
organised and funded by Artist

The relationship between control and computation has been in the focus of the research of Paul Caspi, leading, among other things, to the development of the Lustre language which combines the high-level description of control loops as viewed by the control engineer, with insights coming from the theory and practice of programming languages.

http://www.artist-embedded.org/artist/Between-Control-and-Software.html

Upcoming Major Workshops and Seminars

within EmSoft / ES Week
Foundations of Component-based Design

September 30th, 2007
Salzburg,Austria
Organised,funded by Artist

The workshop aims to gather together researchers from computer science and electrical engineering and will seek a synthesis between the the underlying paradigms and techniques. The focus is not only on fundamental results but also on their implementation in methods and tools and their concrete application in areas such as automotive, avionics, consumer electronics and automation.

The workshop will address specific challenges such as:

- Foundations and Expressiveness of System Description Formalisms:
  - *basic concepts,*
  - *componentinteraction,*
  - *resource modeling (energy, memory, time, ...),*
  - *combining synchrony vs. asynchrony, event-triggered/data-triggered/ time triggered, separation of concerns;*
- Component-based Design, Methods and Tools:
  - *analysis methods (compositional verification techniques; resource usage);*
  - *designmethods(property preserving structuring principles; refinement/ implementation relations)*
  - *tradeoffs between predictability and efficiency*
  - *implementation methodologies and tools*
- Application Scenarios and Relevant Case Studies

http://www.artist-embedded.org/artist/-Foundations-of-Component-based-.html

WESE'07
Workshop on Embedded Systems Education

October 4-5, 2007
Salzburg,Austria
Organised,funded by Artist

It is widely recognized that the embedded system domain is a multidisciplinary one, requiring a large variety of skills from control and signal processing theory, electronics, computer engineering and science, telecommunication, etc., as well as application domain knowledge.

This has motivated a recent but ever growing interest in the question of educating specialists in this domain and this has also been recognized as a particularly difficult problem.

This third workshop on the subject aims to bring researchers, educators, and industrial representatives together to assess needs and share design, research, and experiences in embedded systems education.

Organisers

- Jeff Jackson
  University of Alabama, USA
- Martin Törngren
  Royal Institute of Technology, Sweden

Programme Committee

- Reiner Hartenstein, Kaiserslautern University of Technology, Germany
- Yann-Hang Lee, Arizona State University, USA
- Jogesh Muppala, The Hong Kong University of Science and Technology, Hong Kong
- Kenneth G. Ricks, The University of Alabama, USA
- Falk Salewski, Aachen University, Germany
- Chi-Sheng (Daniel) Shih, National Taiwan University
- Stewart Tansley, Microsoft, Redmond, WA, USA
- Wayne Wolf, Princeton University, USA

http://www.artist-embedded.org/artist/WESE-07.html

Integrated Modular Avionics
ARTIST2 meeting on Integrated Modular Avionics

November 12-13, 2007
Rome, Italy
Organised, funded by Artist

Today, the exponentially increasing diversity of airborne systems results in an ever increasingnumberofcomputers and controllers for system management, monitoring, and control. The development of specific ad-hoc solutions causes increases in costs, which in turn impacts purchase prices and operational costs. To overcome this, standardization principles and reuse of function units are now considered, via Integrated Modular Avionics.

Integrated Modular Avionics (IMA) has set the principles of standardized components and interfaces of hardware and software in aircraft. These principles have been applied for the first time in the development of the Airbus A380. Further developing IMA raises a number of issues that require fundamental research efforts, in tight coordination with engineering needs.

ARTIST2, the European Network of Excellence on embedded systems has decided to organize, as part of its activity on «scientific challenges in specific industrial sectors», a two-day workshop dedicated to Systems, Software, and Architecture, aspects of IMA.

http://www.artist-embedded.org/artist/Integrated-Modular-Avionics.html

Other Conferences

The full set of conferences referenced by Artist is available here:
http://www.artist-embedded.org/artist/Conferences-63.html

See Also

FORMATS'07
October 3-5, 2007
Salzburg, Austria
organised with Artist partners
Researchers interested in semantics, verification and performance analysis study models such as timed automata and timed Petri nets. the digital design community focus on propagation and switching delays while designers of embedded controllers have to take account for the time taken by controllers to compute their responses after sampling the environment.
http://www.ulb.ac.be/di/formats07/

RTSS 2007
December 3-6, 2007
Tucson, Arizona, USA
organised with Artist partners
RTSS provides a forum for the presentation of high-quality, original research covering all aspects of real-time systems design, analysis, implementation, evaluation, and case-studies. RTSS'07 continues the trend of making RTSS an expansive and inclusive symposium, looking to embrace new and emerging areas of real-time systems research.
http://www.rss.org/

Embedded Systems Week 2007

Sept 30th - Oct 5th, 2007

HARDWARE - SOFTWARE CODESIGN AND SYSTEM SYNTHESIS

CODES+ISSS

The Conference covers all aspects of embedded computing systems, including but not limited to:

- *High-level, architectural and system-level synthesis* Synthesis, partitioning, design space exploration
- *HW/SW co-design* Co-design methodologies, HW/SW interface.
- *Spec languages and models* System-level models & semantics, timing analysis, power, formal properties, heter. systems.
- *Simulation and verification* HW/SW cosimulation, HW acceleration, test meth., design for testability
- *Power-aware design meth.* Power mgt / modeling, low-power design meth.
- *ES architecture* Application-specific architectures, memory & comm. optimization
- *Multiprocessors and NoC* Multiproc architectures, communication protocols, design space exploration, MPSoC and NoC.
- *Embedded software* Compilers, virtual machines, scheduling, power-aware OS, RT support and middleware.
- *Application-specific design and algorithms* Network & media processors, hardware accelerators, reconfigurable processors, securities
- *Industrial practices and design case studies* New approaches to areas such as cell phones, sensor networks, automotive, multimedia, med. systems.
- *Emerging techniques* arising from increased heterogeneity, new technologies / applications, Networked embedded systems.

ANNUAL ACM SIGBED CONFERENCE ON EMBEDDED SYSTEMS SOFTWARE

EMSOFT ***Sponsored by Artist2***

EMSOFT is an annual ACM Conference on Embedded Systems Software sponsored by ACM SIGBED (Special Interest Group on Embedded Systems).

Embedded software must meet demanding criteria for correctness, performance, power consumption, and development cost. EMSOFT aims at covering all aspects of embedded software with focus on principles of embedded software development. Topics of Interest include (but are not limited to):

- Design and implementation of embedded software
- Modeling and validation
- Model- and component-based software design and analysis
- Programming languages and compilers
- Software engineering and programming methodologies
- Scheduling and execution time analysis
- Operating systems and middleware
- QoS management and performance analysis
- Hardware-dependent software and interfaces
- Networked embedded systems and security
- Software for distributed/ multiprocessor embedded systems
- Application areas, e.g. automotive, avionics, telecommunication, and multimedia

COMPILERS, ARCHITECTURE, AND SYNTHESIS FOR EMBEDDED SYSTEMS

CASES

Conference topics include, but are not limited to, the following areas:

- Application-specific and domain-specific embedded systems
- Compilation techniques that focus on embedded architectures
- Dynamic compilation and managed runtime environments for embedded systems
- Design, specification, and synthesis of embedded systems
- Customizable processors and digital signal processors
- Embedded uses of instruction-level parallelism, including VLIW, EPIC and superscalar
- Embedded system integration and testing
- Multiprocessing on chip (hardware and software issues)
- Memory management, smart caches and compiler controlled memories
- Novel architectures and micro-architectures for embedded systems
- Low-power architectures and compilation, power vs. performance tradeoffs
- Profiling, measurement, and analysis techniques of embedded applications
- Reconfigurable embedded computing systems
- Validation, verification, and debugging techniques for embedded software
- VLSI and circuit techniques for embedded system design

http://www.esweek.org/

Component-based Design and Implementation of Avionics Systems

Olaf Heinzinger and Maria Sorea
EADS Innovation Works, Germany

The development of modern aircraft and avionics systems has reached a degree of complexity where a highly integrated design process between various disciplines, such as electrical engineering, computer science, and control theory is necessary.

Sharc, an unmanned aerial vehicle (UAV) is an example of such an avionics system developed at EADS. With a maximum take-off weight of 190 kg, the system can accommodate 50 kg of mission equipment in its payload compartments. Sharc is equipped with a digital flight-control unit, a laser altimeter, and control and data link. Sharc has been designed as an unmanned aerial vehicle without hydraulic compo-nents, the rotors being controlled by means of electrical actuators.

Autonomous systems such as Sharc are mostly computer-controlled and are character-ized by real-time computations, as well as strong interactions and coupling effects between the components and the environment. These aspects, together with require-ments of fault tolerance, make the design, analysis and verification of such embedded systems inherently difficult.

In order to design such complex systems, and to gain confidence in the validity of the design itself, precise models of the components and of the surrounding relevant envi-ronment need to be constructed. As shown in the figure on the right, the system model of Sharc consists not only of the aircraft components but also of atmospheric phenomena and terrain representation. We use mainly Stateflow and Simulink for expressing these models. The models are validated using a simulator or rapid prototype, allowing also for exploration against various "what-if" scenarios. To attain the desired – or required – level of confidence in correct behaviour, however, mere testing and simulation is usually insufficient; in fact, it is argued that the kind of reliability required for highly safety-critical applications cannot be achieved without a careful formal analysis of the mechanisms and algorithms involved. Besides the usual simulation runs we are exploring state-of-the-practice verification methods, such as model checking for hybrid systems to achieve a high level of confidence into the functional behaviour of the design. Model-checking is used for increasing the ro-bustness of the design using a technique called model-based safety analysis. The models are extended with "what-if" scenarios for modeling both failure of the aircraft and unexpected behaviour of the environment, and model checking is used to automatically examine all possible failure scenarios. We are currently investigating the possibility of using model-checking for exhaustively exploring sufficiently precise abstractions of the aircraft and environmental models. The main challenge here is to develop sufficiently powerful verification techniques for distributed, real-time, embedded systems with non-linear dynamics. In addition, for industrial dissemination, these verification techniques need to be fully integrated into the design cycle and the development tool chain.

Following design validation, the subsystem models are used for analyzing the interac-tion of the represented components with the developed avionics system running on its target hardware (hardware-in-the-loop simulation). This reduces the risk when operat-ing and testing the operational prototype. For the creation of such a development envi-ronment, a distributed set of embedded systems is the best solution for a flexible test and development process and for a realistic system representation including bus and communication structures of an integrated modular avionics (IMA) architecture.

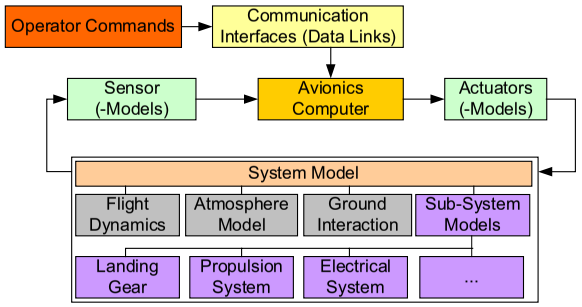

Performance / Power Measurements for Embedded Multi-Core Processors

Dr. Jörg Brakensiek
Nokia Research Center, SRC Bochum

Recently main PC computing performance gains has been achieved by increasing the number of CPUs instead of higher clock rates because of technological limitations. In the embedded world other restrictions exist. Nevertheless, multi core solutions might also be an option here. From the embedded perspective, the processing efficiency (Eq. 1) is more important than pure performance numbers.

$$\text{Efficiency} = \frac{\text{ComputingPerformance}}{\text{Energy}} \qquad \text{(Eq. 1)}$$

So the question needs to be answered, whether a multi-core solution is actually more efficient than single core one (Efficiency Gain > 1, Eq. 2).

$$\text{EfficiencyGain} = \frac{\text{Efficiency}_{MultiCPU}}{\text{Efficiency}_{SingleCPU}} \qquad \text{(Eq. 2)}$$

The Nokia Research Center in Bochum, Germany has been investigating the embedded ARM multi-core processor test chip (ARM CT11 MPCore [1]) on the ARM Versatile EB [2] in order to find specific answers to the above question.

We have been analyzing a set of low-level and high-level benchmarks. In order to get comparable results, we used the following Frequency & Number-of-CPUs settings.

$$f_i \cdot N_i = f_j \cdot N_j$$

The measurement results from Dhrystone benchmarks are shown in figures below. They first show the expected linear performance over frequency as well as over number of CPUs (figure 1).

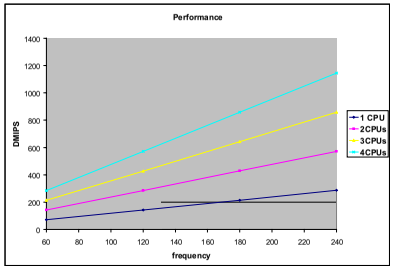

figure 1

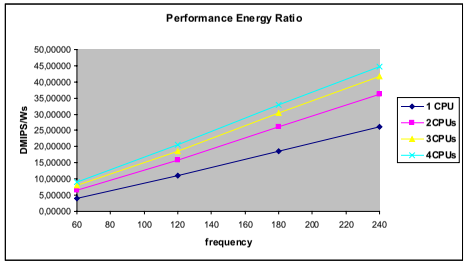

figure 2

When it comes to the performance efficiency (figure 2), the results are showing a linear dependency from the frequency and a non-linear dependency from the number of cores.

In all cases, increasing the number of cores and frequency is increasing the efficiency of the multi-core system.

An example for high-level (i.e. application level) benchmark results is shown in the following figures. The efficiency results and corresponding efficiency gains for different configurations for the multi-threaded FFPlay (an open source media player) is shown. It has to be noticed that the video output has been disabled to eliminate the influence of the low IO bandwidth of the Versatile EB.

FFPlay 25fps Video	1 Thread	2 Thread	4 Thread
	1 CPU @ 240 MHz	1 CPU @ 240 MHz	1 CPU @ 240 MHz
2 CPU @ 120 MHz	1,32	1,29	1,33
4 CPU @ 60 MHz	0,74	1,44	1,40
	2 CPU @ 120 MHz	2 CPU @ 120 MHz	2 CPU @ 120 MHz
4 CPU @ 60 MHz	0,56	1,12	1,05
	2 CPU @ 240 MHz	2 CPU @ 240 MHz	2 CPU @ 240 MHz
4 CPU @ 120 MHz	1,29	1,28	1,26

figure 3

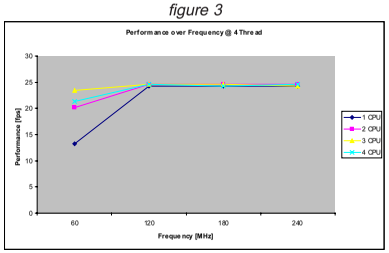

figure 4

It can be seen that the target frame rate (25 fps) can be achieved for all frequency settings above 120 MHz in the 4-threaded case (left figure). Achievable efficiency gains (right figure) are between 20% - 30% (the efficiency value from the first column settings are divided by the one of the 2nd to 4th column). The values in the second and third row are not completely comparable, as the final frame rate cannot be achieved (figure 4).

The current state of the analysis has shown that in case of well distributable workload the efficiency of such a multi-core environment is in many cases much higher compared to a single-core environment.

The analysis has turned out that there is complex dependency of the achievable efficiency from the system configuration and setup, which still requires more understanding with respect to the software and hardware side in order to take full advantage of this technology.

References

[1] Core Tile for ARM11 MPCore, User Guide; ARM Limited, 2005; http://www.arm.com/pdfs/DUI0318C_core_tile_11mpcore_ug.pdf
[2] Versatile EB Board; http://www.arm.com/products/DevTools/EB.html
[3] H. Bothe, "Increase Performance with Embedded Multi-Core Processors", ARM IQ, Volume 5, No 4, 2006.
[4] J. Brakensiek, "Multi Execution Environments @ Multi Cores Use Cases, Requirements and Concepts", Presentation on MultiCore Expo, Santa Clara, 29.03.2007.

Current and Recent Schools and Courses

Artist2 / UNU-IIST School in China - 2007
August 1-10th
Suzhou, China
http://www.artist-embedded.org/artist/-Artist2-UNU-IIST-School-in-China-.html
Organised by the Artist2 NoE

ARTIST2 has organized, in collaboration with UNU-IIST, the 2nd edition of a school on embedded systems design in Suzhou (near Shanghai).

Topics Covered included: MPSoCs, Interaction between Control and Scheduling, Model-based development, Automata and hybrid systems, Semantics Issues of Stateflow, Mixed Continuous and Discrete-event Systems, Verification Techniques, Optimal Scheduling, Controller Synthesis.

Lecturers

- Prof. Karl-Erik Arzen
  Lund University - Sweden
- Prof. Dr. Luca Benini
  University of Bologna - Italy
- Paul Caspi
  Verimag Lab - France
- Prof. Kim Larsen
  Aalborg Univ. - Denmark

Joint EU-China Organisation

The ARTIST2 / UNU-IIST / China Spring School was initialized and organized
jointly by the ARTIST2 Network of Excellence, and UNU-IIST, Macao.

First European-SouthAmerican School for Embedded Systems
August 21-24, 2007
Buenos Aires - Argentina
http://www.artist-embedded.org/artist/-First-European-SouthAmerican-.html
Organised by the Artist2 NoE

The purpose of the school is to foster dynamic research cooperation between groups in Europe and South America, by allowing South-American students (specially graduate), to meet european researchers. We strongly believe this will offer an excellent opportunity to strengthen the relationships with mutual benefit.

The school will be a repeated event on a yearly basis. Besides the lectures given by european researchers, there will be invited talks by southamerican researchers and space (poster session) for graduate students to present and discuss their work.

Topics

- Component-based modeling of heterogeneous real-time systems
- Adaptive Real-time systems
- Networks for embedded control systems

Lecturers

- Prof. Luis Almeida
  Univ. of Aviero - Portugal
- Prof. Gerhard Fohler
  TU Kaiserslautern - Germany
- Joseph Sifakis
  Verimag Lab - France

FOSAD'07
Foundations of Security Analysis and Design
September 9-15, 2007
Bertinoro - Italy

http://www.sti.uniurb.it/events/fosad/
Sponsored by the Artist2 NoE

Security in computer systems and networks emerged as one of the most challenging research areas. The International School on Foundations of Security Analysis and Design (FOSAD) has been one of the foremost events established with the goal of disseminating knowledge in this critical area. The main aim of the FOSAD school is to offer a good spectrum of current research in foundations of security - ranging from programming languages to analysis of protocols, from cryptographic algorithms to access control policies and trust management - that can be of help for graduate students and young researchers from academia or industry that intend to approach the field.

Courses

- API Security and Security Economics
- Low-level Software Security
- Application of Formal Methods to Cryptographic Protocol Analysis
- Trusted Mobile Platforms
- Language-Based Security
- Cryptographic Algorithm Engineering and Provable Security
- Embedded Systems Security and Cryptographic Coprocessors
- Quantitative Aspects in the Analysis of Crypographic Protocols

Applied Software Verification
LASER Summer School on Software Engineering
September 9-15, 2007
Elba - Italy

http://se.inf.ethz.ch/laser/2007/

The LASER school is intended both for researchers (including PhD students) and for professional software engineers and managers who want to benefit from the best in software technology advances. The focus of LASER is resolutely applied, although theory is welcome to establish solid foundations. The format of the school favors extensive interaction between participants and speakers.

The 2007 LASER school is part of the ongoing «Grand Challenge» on software verification, initiated by Tony Hoare. It has a special focus on tools for software verification. This means in particular that it has a highly practical character and will provide participants with a clear view of techologies and tools available today to verify software.

EPSD 2007
Embedded Programmable System Design
September 9-15, 2007
EPFL, Lausanne, Switzerland
http://www.mead.ch/htm/ch/EPSD-Program.html

- Introduction to Embedded Systems
- Software Design Principles
- Real-Time Scheduling and Performance Estimation
- VLIW Architectures
- Embedded Memory Systems
- Memory Architecture Aware Compilation
- Retargetable Compilers
- Automatic Processor Specialization
- From Algorithms to Architectures: A Case Study
- Power Analysis and Low Power Design
- System-Level Power Optimization

The Role of Networking in Embedded Systems

Luis Almeida, University of Aveiro, Portugal
Eduardo Tovar, Polytechnic Institute of Porto, Portugal

Looking at the current scenario in embedded systems we see a permanently growing role of networking, ranging from the connection of autonomous devices such as cell phones, PDAs, laptops and their peripherals, to the provision of pervasive access to multimedia and telecommunication networks, to the deployment and operation of large-scale sensor networks, to intelligence distribution in complex embedded systems, or even, at a small physical scale, to connect multiple processing cores within Systems-on-Chip (SoCs).

In this vast horizon, the ongoing Artist2 NoE's activity on Dynamic and Pervasive Networks of the ARTIST2 Adaptive Real-Time Systems (ART) cluster focuses on Wireless Sensor Networks (WSNs), Mobile Ad-hoc Networks (MANETs) and Networked Embedded Systems (NESs), areas in which many open challenges remain, and many new problems arose in the past few years. For example, energy-aware communication is turning out to become a major research challenge for WSNs, imposing innovative and efficient networking protocols that manage communications periodicity, nodes synchronization and transmitting power; As WSNs grow into very large-scale networks with thousands of nodes and more, efficient data aggregation becomes essential being imperative that its time-complexity does not depend on the number of sensor nodes; Quality-of-Service (QoS) adaptation and the collaborative computing paradigms require protocol mechanisms to monitor instantaneous bandwidth usage, enforce minimum agreed QoS levels (e.g. through contracts and traffic policing) and leverage the access to free bandwidth (to increase QoS whenever possible); Higher software integration in distributed embedded systems requiring integrated global resource management together with effective and efficient temporal partitioning (e.g., using hierarchical scheduling techniques), as well as flexible mapping between software and hardware architectures; Replacement and/or extension of wired with wireless networking technologies, coping with more error-prone channels and security risks but profiting from simplified deployment and elimination of cabling.

Moreover, distributed sensing, actuation and cooperative computing involving small and tiny computing platforms appear as a basilar functionality in an ever crescent range of applications, including surveillance, environment and critical infrastructures monitoring, disaster recovery operations, distributed control, military operations, etc. The requirements imposed by these diverse applications necessarily imply different trade-off options on supported functionality, quality of service, efficiency, platforms, protocols, architectures, etc.

Among the initiatives being carried out in this activity we highlight two, the real-time communication and collaboration in wireless networks and cyber-physical systems lead by the Politechnic Institute of Porto, and the network support for dynamic reconfiguration, which was the topic of the NeRES workshop that took place on the 2nd of April, this year, in Aveiro, Portugal (http://www.artist-embedded.org/artist/-NERES-2007-.html). These two initiatives are briefly described in the following articles.

Adaptive Real Time for Embedded Systems

Giorgio Buttazzo, Scuola Superiore Sant'Ana, Pisa, Italy

Achieving adaptivity in embedded real-time systems is a complex task that requires expertize from several disciplines, including operating systems, scheduling theory, network communication, control theory, and quality of service management. To cover these issues, the ART cluster includes the following activities:

1. A common infrastructure for adaptive real-time systems show how current operating systems and network protocols can be extended to support emerging real-time applications that exhibit a high degree of complexity and operate in dynamic environments.
2. Flexible Scheduling and Resource Management provides models, policies and analysis techniques to efficiently manage the available resources.
3. Dynamic and Pervasive Networking addresses the numerous research challenges in the frameworks of Wireless Sensor Networks, Mobile Ad-hoc Networks and Embedded Networked Systems.

Networks for Reconfigurable ES

Luis Almeida, University of Aveiro, Portugal
Paulo Pedreiras, University of Aveiro, Portugal

Reconfigurability has long been recognized as a way to improve efficiency in the use of system resources, for example, when a system undergoes variable load situations, when it evolves during its lifetime or even when faults affect part of its structure. This means that reconfigurability, in a broad sense, may be beneficial to areas that range from Quality of Service (QoS), e.g., when the number of system users or the environmental operating conditions vary, to Dependability, e.g., through graceful degradation.

However, achieving reconfigurability may conflict with operational goals such as continued real-time and safe operation, and it becomes more difficult when the system is distributed, requiring adequate support from the network. One approach that has been followed in certain application domains to cope with such difficulty is the use of multiple operational modes, which are statically defined off-line. Nevertheless, more flexible approaches to reconfigurability are needed to improve resource efficiency in a vast range of applications, exploiting paradigms such as flexible modes, flexible scheduling, dynamic QoS management, stateful schedules, etc, particularly at the network level. This opens the way to keep costs low while improving dependability as the overall system complexity increases.

The NeRES workshop (Networks for Reconfigurable Embedded Systems) that took place on April 2nd in Aveiro, Portugal, aimed at discussing these problems associated to dynamic reconfiguration in distributed embedded systems and, particularly, the role of the network to support it. It gathered 26 participants from 15 institutions in 6 countries, with one industrial representative and several other academic participants presenting industrial case studies. There were 13 presentations covering aspects that ranged from flexible middleware, namely based on components, on resource contracts, on services and on the support for flexible scheduling, to dependability, integration, wireless mobile ad-hoc communication, intelligent telecommunication networks, industrial automation, automatic control systems, automotive and avionic systems.

More information can be found here:
http://www.artist-embedded.org/artist/-NERES-2007-.html

Focus on Adaptive Real Time

ARTIST Newsletter n°4
September 2007

Real Time Communication / Collaboration in Wireless Networks and Cyber-Physical Systems

Björn Andersson, Polytechnic Institute of Porto, Portugal
Mário Alves, Polytechnic Institute of Porto, Portugal

The use of wireless communication networks has undergone a revolution during recent years, in application areas such as factory automation, home automation, vehicle-to-vehicle communications and wireless sensor networks (WSN). Several communication standards are established and transceivers are commercially available at low cost. Unfortunately, they were not designed for satisfying real-time requirements. Researchers of the ARTIST2 NoE's Adaptive Real-Time Systems (ART) cluster are addressing these issues from the complementary perspectives of: (i) the use of commercially available technologies versus the use of new solutions untethered by existing standards; (ii) the use of time-triggered paradigms versus event-triggered paradigms; (iii) the provision of short-term solutions recognizing the needs of companies for wireless solutions with mature implementations and compliant with standards in order to simplify interoperability versus the need to push the state-of-art and explore new innovative solutions. ART-WiSe and WiDOM are two representative ongoing research frameworks addressing these complementary perspectives.

The ART-WiSE (Architecture for Real-Time communications in Wireless Sensor networks, http://www.hurray.isep.ipp.pt/art-wise) research framework aims at the specification of a scalable two-tiered communication architecture for improving the timing and reliability behaviour of WSNs. One of the major goals is to use, as far as possible, existing standard communication protocols and commercial-off-the-shell (COTS) technologies – IEEE 802.15.4/ZigBee for Tier 1 and IEEE 802.11 for Tier 2 (Fig. 1).

Results obtained thus far include the provisions of methodologies to analyse and dimension star and cluster-tree 802.15.4/ZigBee networks, namely being able to compute throughput and message delay bounds for the Guaranteed Time Slot (GTS) mechanism and ZigBee Router's buffer requirements in cluster-tree networks. Important add-ons to these protocols that are backward compatible, have already been proposed and tested: (i) a traffic differentiation mechanism for CSMA/CA to provide more guarantees to high priority messages by appropriate tuning of MAC parameters; (ii) an implicit GTS allocation mechanism (i-GAME) allowing improved bandwidth utilization and adaptation by nodes sharing a GTS; (iii) beacon/superframe scheduling in ZigBee cluster-tree networks enabling a synchronized cluster-tree WSN where each cluster may operate with different and low duty-cycle, thus prolonging network lifetime.

An open-source toolset for the IEEE 802.15.4/ZigBee protocols have been made publicly available: http://www.open-zb-net, which includes: (i) the implementation of the IEEE 802.15.4 protocol in TinyOS, for both the MICAz and TelosB motes; (ii) the implementation of the ZigBee Network Layer for supporting synchronized multiple cluster topologies (the Cluster-Tree topology) in TinyOS, for the TelosB motes; (iii) a simulation model of the IEEE 802.15.4 protocol in OPNET; (iv) tools for timing analysis and network dimensioning.

The Artist2 NoE's ART cluster has also recognized that no existing communication standard performs well for sporadic message streams with real-time requirements. For this reason, a novel MAC protocol, dubbed Wireless Dominance Protocol (WiDOM), was designed for wireless systems (http://www.hurray.isep.ipp.pt/activities/WiDOM/). This protocol gives the wireless channel a similar behavior as a Controller Area Network (CAN) bus. It is prioritized and this can be achieved even without having the ability to listen and transmit simultaneously. Because of the prioritization, it is possible to compute message response-times of sporadic message streams.

The WiDOM protocol elects the computer node with the highest priority (lowest number) and gives it access to the medium. This election procedure can also be used to compute the minimum value of sensor readings distributed on different computer nodes and, remarkably, this computation can be performed with a time-complexity that is independent of the number of computer nodes. This procedure forms a building block for other useful calculations; for example, it is possible to efficiently extract an interpolation of sensor readings and this can be performed with a time-complexity that is independent of the number of computer nodes. This is a crucial asset for addressing problems in future Large-Scale Dense Sensor Networks for Cyber-Physical Systems.

Figure 1:
Example of the ART-WiSe two-tiered network architecture

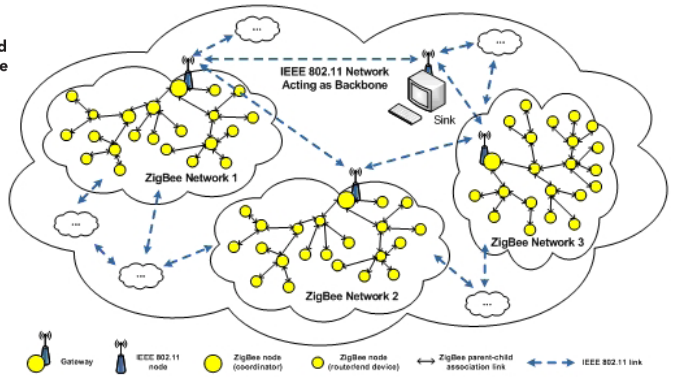

6. Education

Embedded software and systems are at the intersection of electrical engineering, computer engineering, and computer science, with, increasing importance, in mechanical engineering. Despite the clear need for knowledge of systems modeling and analysis (covered in electrical and other engineering disciplines) and analysis of computational processes (covered in computer science), few academic programs have integrated these disciplines into a cohesive program of study.

This will continue and be reinforced in the ArtistDesign NoE.

6.1 Educational Methods for Embedded Systems Design

Artist2 and the preceding Artist FP5 project have a continuing tradition for organising events and published materials to further Education in Embedded Systems Design.

The workshops described here address educational issues for embedded systems in general.

Workshop on Embedded Systems Education (WESE 2008)
October 23rd, 2008 Atlanta, Georgia - USA (within ESWEEK)
http://www.artist-embedded.org/artist/-WESE-08-WS-on-Embedded-Systems-.html

Particular topics of interest include but are not limited to:

- Industrial needs regarding embedded systems education
- Embedded systems curricular design and implementation
- Control and signal processing issues
- Computer science issues
- Real-time computing issues
- Distributed systems issues
- Architecture and design issues
- Hardware/software co-design
- Hands-on experiences and labs
- Teaching embedded systems

Formally, the 2008 edition of WESE is part of rthe ArtistDesign NoE.

Workshop on Embedded Systems Education (WESE 2007)
October 4-5, 2007 Salzburg, Austria (within ES Week)
http://www.artist-embedded.org/artist/-WESE-07-.html

It is widely recognized that the embedded system domain is a multidisciplinary one, requiring a large variety of skills from control and signal processing theory, electronics, computer engineering and science, telecommunication, etc., as well as application domain knowledge.

This has motivated a recent but ever growing interest in the question of educating specialists in this domain and this has also been recognized as a particularly difficult problem.

This third workshop on the subject aims to bring researchers, educators, and industrial representatives together to assess needs and share design, research, and experiences in embedded systems education. Industrial needs regarding embedded systems education

Workshop on Embedded Systems Education (WESE 2006)
October 26th, 2006 Seoul, Korea
http://www.artist-embedded.org/artist/-EmSoft-06-Workshop-on-Embedded-.html

Workshop on Embedded Systems Education (WESE 2005)
September 22nd, 2005 Jersey City, New Jersey - USA
http://www.artist-embedded.org/artist/-ARTIST2-Workshop-Workshop-on-.html

International Collaboration Day on Education
October 11th 2003 – Philadelphia
http://www.artist-embedded.org/artist/International-Collaboration.html

Artist-FP5/ACM Transactions Embedded Computing Systems Special Issue on Education *(publication)*
http://www.artist-embedded.org/artist/ACM-Special-Issue-on-Education.html

Embedded systems applications now include a very large proportion of the advanced products designed in the world, spanning transport (avionics, space, automotive, trains), electrical and electronic appliances (cameras, toys, television, washers, dryers, audio systems, and cellular phones), process control (energy production and distribution, factory automation and optimization), telecommunications (satellites, mobile phones and telecom networks), and security (e-commerce, smart cards), etc.. The relative weight of software in the value of embedded systems is constantly expanding. The extensive and increasing use of embedded systems and their integration in everyday products marks a significant evolution in information science and technology.

There is now a strategic shift in emphasis for embedded systems designers: from simply achieving feasibility, to achieving optimality. Optimal design of embedded systems means targeting a given market segment at the lowest cost and delivery time possible. Optimality means seamless integration with the physical and electronic environment while respecting real-world constraints such as hard deadlines, reliability, availability, robustness, power consumption, and cost. In our view, optimality can only be achieved through the emergence of embedded systems as a discipline in its own right.

An important factor for the emergence of embedded systems as a discipline is the existence of integrated curricula for training engineers and researchers, able to tackle a range of topics which until now had been spread across many different areas, including: general computer science and engineering, real-time computing, systems architecture, control and signal processing, security and privacy, networking, mathematics, electronics.

This special issue of the ACM Transactions in Embedded Computing Systems aims to provide the basis for integrated undergraduate and graduate curricula covering the essential areas of knowledge for tomorrow’s embedded systems engineers and researchers.

Artist FP5 Guidelines for a Graduate Curriculum on Embedded *(publication)*
http://www.artist-embedded.org//docs/Publications/Education.pdf

The design of real-time embedded systems requires skills from three specific disciplines: control theory, computer science, and electronic engineering, and their combination. This often involves experts from differing backgrounds, who do not recognize that they address different issues from complementary angles.

The motivation for defining a specific curriculum in embedded systems is mainly to promote the training of engineers with expertise in the above three disciplines.

6.2 Summer Schools Organized and/or sponsored by Artist2

6.2.1 Organised and funded by the Artist2 NoE

ARTIST2 Summer School 2008 in Europe
September 8-12, 2008 Autrans (near Grenoble), France
http://www.artist-embedded.org/artist/-ARTIST2-Summer-School-2008-.html

The Summer School offers a number of foundational tutorials, accompanied by a selection of lectures on exciting emerging technologies and industrial applications - given by leading scientific and/or industrial experts.

ARTIST2 South-American School for Embedded Systems 2008
August 25-29, 2008 Universidade Federal de Santa Catarina, Florianopolis, Brazil
http://www.artist-embedded.org/artist/-ARTIST-2-South-American-School-.html

Second edition of the ARTIST South American School.

Artist2 Summer School in China 2008
July 12-18, 2008 Shanghai, China
http://www.artist-embedded.org/artist/-Artist2-Summer-School-in-China-.html

ARTIST2 will organize the 3rd edition of a school on Embedded Systems Design in Shanghai. This year, the school is organized in collaboration with the SEI/ECNU and the LIAMA.

Real-Time Kernels for Microcontrollers: Theory and Practice
June 23-25, 2008 Pisa, Italy
http://www.artist-embedded.org/artist/-Real-Time-Kernels-for-.html

The course on Real-Time Kernels for Microcontrollers aims to introduce the basic concepts of Real-time Systems targeted to Embedded Systems, which are often implemented using microcontrollers. The course will briefly illustrate the theoretical background of real-time scheduling, resource-aware techniques, and wireless communication based upon the IEEE 802.15.4 protocol.

ARTIST2 Graduate Course on: Automated Formal Methods for Embedded Systems 2008
June 16-24, 2008 DTU - Lyngby, Denmark
http://www.artist-embedded.org/artist/-Automated-Formal-Methods-for-.html

In the lectures, we will introduce a comprehensive set of state-based models as well as automatic procedures for their analysis. The exercise classes will complement this by providing hands-on experience with appropriate verification tools.

ARTIST2 Graduate Course on Embedded Control Systems
May 26-30, 2008 Stockholm, Sweden
http://www.artist-embedded.org/artist/-Graduate-Course-on-Embedded-.html

The course provides an account of state of the art theory and techniques that address the connection and integration of the areas of Control systems and Embedded systems.

First European-SouthAmerican School for Embedded Systems
August 21-24, 2007 Universidad Argentina de la Empresa (UADE), Buenos Aires – Argentina
http://www.artist-embedded.org/artist/-First-European-SouthAmerican-.html

The purpose of the school is to foster the well established and dynamic research cooperations in the field of embedded systems between groups in Europe and South America, by allowing south-american students (specially graduate), to meet european researchers.

Artist2 / UNU-IIST School in China – 2007
August 1-10, 2007 Suzhou (near Shanghai), China
http://www.artist-embedded.org/artist/-Artist2-UNU-IIST-School-in-China-.html

ARTIST2 will organize, in collaboration with UNU-IIST, the 2nd edition of a school on embedded systems design in Suzhou (near Shanghai).

ARTIST2 PhD Course on: Automated Formal Methods for Embedded Systems
June 4-12, 2007 DTU - Lyngby, Denmark
http://www.artist-embedded.org/artist/-ARTIST2-PhD-Course-on-Automated,851-.html

In the lectures, we will introduce a comprehensive set of state-based models as well as automatic procedures for their analysis. The exercise classes will complement this by providing hands-on experience with appropriate verification tools.

ARTIST2 Graduate Course on Embedded Control Systems
May 7-11, 2007 Lund, Sweden
http://www.artist-embedded.org/artist/-ARTIST-Graduate-Course-on-Embedded-.html

The objective of the course is to provide an overview of the main principles and technologies for supporting the development of embedded control systems.

Real-Time Microcontroller Systems: OSEK Standard and experiments on µcontroller devices
March 26-28, 2007 RETIS Laboratory, Scuola Superiore Sant'Anna, Pisa, Italy
http://www.artist-embedded.org/artist/-OSEK-Standard-and-Multicore-.html

Training course on Real-Time Systems for Microcontrollers: OSEK Standard and experiments on microcontroller devices *Organised in conjunction with Evidence Srl*

ARTIST2 - MOTIVES 2007
February 19-23, 2007 Trento, Italy
http://www.artist-embedded.org/artist/-MOTIVES-2007-.html

ARTIST2 Winter School 2007 offers foundational tutorials and lectures on exciting emerging technologies and industrial applications - given by leading scientific and industrial experts.

First European Laboratory on Real-Time and Control for Embedded Systems
July 10-14, 2006 Pisa, Italy
http://www.artist-embedded.org/artist/-First-European-Laboratory-on-Real-.html

Real-Time distributed embedded systems play a crucial role in our society including several application domains such as automotive, telecommunications, robotics, and multimedia systems. These systems generally work under precise timing constraints, to achieve the required level of performance and predictability. Consequently, embedded systems design requires expertise in several disciplines, including control theory, networking, real-time computing, and operating systems.

ARTIST2 / UNU-IIST Spring School in China 2006
April 3-15, 2006 Xi'an, China
http://www.artist-embedded.org/artist/-ARTIST2-UNU-IIST-China-School-.html

The first ARTIST / UNU-IIST Spring School gathered more than 50 participants, of which approximately 40 were students from the top universities in mainland China.

ARTIST2 Graduate Course on Embedded Control Systems
April 3-7, 2006 Prague, Czech Repbulic
http://www.artist-embedded.org/artist/-ARTIST2-Graduate-Course-on-.html

The objective of the Course is to provide an overview of the main principles and technologies for supporting the development of embedded control systems.

ARTIST2 Summer School 2005
September 29th - October 2nd 2005 Nässlingen, Sweden
http://www.artist-embedded.org/artist/-ARTIST2-Summer-School-2005-.html

ARTIST2 Summer School on Component & Modelling, Testing & Verification, and Statical Analysis of Embedded Systems.

6.2.2 Sponsored by the Artist2 NoE or Organised with Artist partners

ARTIST2 has provided support for selected schools and seminars:

RNTS'08 *organised with Artist partners*
http://www.artist-embedded.org/artist/-RNTS-08-.html *October 16th, 2008*
16th International Conference on Real-Time and Network Systems Rennes, France.

FCC'08 *sponsored by Artist*
http://www.artist-embedded.org/artist/-FCC-08-.html *June 26th, 2008*
This workshop focused on approaches that combine and relate symbolic and computational protocol analysis. Over the last few years, there has been a spate of research results in this area. One set of results establish correspondence theorems between the two models, in effect showing that for a certain class of protocols and properties, security in the symbolic model implies security in the computational model. In other work, researchers use language-based techniques such as process calculi and protocol logics to reason directly about the computational model. Several projects are investigating ways of mechanizing computationally sound proofs of protocols. The workshop seeks results in this area of *computationally sound protocol analysis: foundations and tools*.

MPSoc 2008 *sponsored by Artist* *organised with Artist partners*
http://www.artist-embedded.org/artist/-MPSoc-2008-.html *June 23rd, 2008*
MPSoC is a pluridisciplinary forum bringing together key R&D actors from the different fields required to design heterogeneous multiprocessor SoC (MPSoC).

UML&AADL'2008 *sponsored by Artist* *organised with Artist partners*
http://www.artist-embedded.org/artist/-UML-AADL-2008-.html *April 2nd, 2008*
All aspects of the representation, analysis, and implementation of Distributed, Real-time and Embedded systems (DRE) system behaviour and/or architecture models.

Scopes 2008 *sponsored by Artist* *organised with Artist partners*
http://www.artist-embedded.org/artist/-Scopes-2008-.html *March 13th, 2008*
SCOPES focuses on the software generation process for modern embedded systems. Topics of interest include all aspects of the compilation process, starting with suitable modeling and specification techniques and programming languages for embedded systems. The emphasis of the workshop lies on code generation techniques for embedded processors.

DATE'08 *sponsored by Artist* *organised with Artist partners*
http://www.artist-embedded.org/artist/-DATE-08-.html *March 10th, 2008*
Originally in the area of design automation, the DATE conference and exhibition has developed a very active embedded software track with 13 sessions on software topics and 2 special days on automotive and on dependable systems in 2008.

EmSoft'07 *sponsored by Artist* *organised with Artist partners*
http://www.artist-embedded.org/artist/-EmSoft-2006-Embedded-Software-.html *Oct 1st, 2007*
EMSOFT aims at covering all aspects of embedded software with focus on principles of embedded software development.

Embedded Systems Week 2007 *sponsored by Artist*
http://www.artist-embedded.org/artist/-Embedded-Systems-Week-2007-.html *Sept 30th, 2007*
Embedded Systems Week brings together conferences, tutorials and workshops centered on various aspects of embedded systems research and development.

EPSD 2007 *sponsored by Artist*
http://www.artist-embedded.org/artist/-EPSD-2007-.html *September 10th, 2007*
Advanced engineering courses will be offered by the Swiss Federal Institute of Technology, Lausanne, Switzerland, during summer period 2007.

FOSAD 2007 *sponsored by Artist*
http://www.artist-embedded.org/artist/-FOSAD-2007-7th-International-.html *Sept 9th, 2007*

The International School on Foundations of Security Analysis and Design (FOSAD) has been one of the foremost events established with the goal of disseminating knowledge in this critical area. The main aim of the FOSAD school is to offer a good spectrum of current research in foundations of security - ranging from programming languages to analysis of protocols, from cryptographic algorithms to access control policies and trust management - that can be of help for graduate students and young researchers from academia or industry that intend to approach the field.

UML&AADL'2007 *sponsored by Artist* *organised with Artist partners*
http://www.artist-embedded.org/artist/-UML-AADL-2007-.html *July 14th, 2007*
This workshop seeks contributions from researchers and practitioners interested in all aspects of the representation, analysis, and implementation of DRE behaviour and/or architecture models.

CAV 2007 *sponsored by Artist*
http://www.artist-embedded.org/artist/-CAV-2007-.html *July 3rd, 2007*
CAV'07 is the 19th in a series dedicated to the advancement of the theory and practice of computer-aided formal analysis methods for hardware and software systems.

FMGALS'2007 *sponsored by Artist*
http://www.artist-embedded.org/artist/-FMGALS-2007-.html *May 29th, 2007*
Third International Workshop on Formal Methods for Globally Asynchronous Locally Synchronous Design

SCOPES 2007 *sponsored by Artist*
http://www.artist-embedded.org/artist/-SCOPES-2007-.html *April 20th, 2007*
SCOPES focuses on the software generation process for modern embedded systems. Topics of interest include all aspects of the compilation process, with emphasis on code generation techniques for embedded processors.

HSCC'07 *sponsored by Artist*
http://www.artist-embedded.org/artist/-HSCC-07-.html *April 3rd, 2007*
The conference, tenth in a series of succesful annual meetings, is dedicated to research in embedded reactive systems involving the interplay between symbolic/switching and continuous dynamical behaviors.

SLA++P 2007 *sponsored by Artist*
http://www.artist-embedded.org/artist/-SLA-P-2007-.html *March 31st, 2007*
Model-driven High-level Programming of Embedded Systems (formerly "Synchronous Languages, Applications, and Programming")

ARCS 2007 *sponsored by Artist*
http://www.artist-embedded.org/artist/-ARCS-2007-.html *March 12th, 2007*
ARCS 2007 will cover basic technology, architecture, and application of computing systems. Autonomic or Proactive computing may help to manage the increasing complexity of computing systems.

CASTNESS'07 Workshop and School *sponsored by Artist* *organised with Artist partners*
http://www.artist-embedded.org/artist/-CASTNESS-07-Workshop-and-School,235-.html
January 15th, 2007
Computing Architectures and Software Tools for Numerical Embedded Scalable Systems

CASTNESS'07 Workshop and School *sponsored by Artist*
http://www.artist-embedded.org/artist/-CASTNESS-07-Workshop-and-School-.html
January 15th, 2007
Computing Architectures and Software Tools for Numerical Embedded Scalable Systems

Synchron 2006 *sponsored by Artist*
http://www.artist-embedded.org/artist/-Synchron-2006-.html *November 27th, 2006*
This workshop is devoted to all aspects of synchronous programming: languages, compiling techniques, formal methods, programming environments, execution platforms, semantics issues, code generation.

JTRES 2006 *sponsored by Artist*
http://www.artist-embedded.org/artist/-JTRES-2006-Java-Technologies-for-.html
October 11th, 2006
Real-time and Embedded Java. This workshop seeks to identify remaining challenging problems remaining to be solved, and to report results and experience gained by researchers.

MARTES 2006 *sponsored by Artist*
http://www.artist-embedded.org/artist/-MARTES-2006-.html *October 2nd, 2006*
This workshop gathers researchers and industrial practitioners to survey modeling and model-based analysis of distributed, real-time and embedded systems.

ADSD 2006: Advanced Digital Systems Design *sponsored by Artist*
http://www.artist-embedded.org/artist/-Advanced-Digital-Systems-Design-.html *Sept 25th, 2006*
Design course for multimillion-transistor Systems-on-Chip and other state-of-the-art embedded products. The course spans from purely digital-design topics to some compiler-related issues.

FOSAD 2006: 6th International School on Foundations of Security Analysis and Design
sponsored by Artist
http://www.artist-embedded.org/artist/-FOSAD-2006-6th-International-.html *Sept 10th, 2006*
The International School on Foundations of Security Analysis and Design (FOSAD) has been one of the foremost events established with the goal of disseminating knowledge in this critical area. The main aim of the FOSAD school is to offer a good spectrum of current research in foundations of security - ranging from programming languages to analysis of protocols, from cryptographic algorithms to access control policies and trust management - that can be of help for graduate students and young researchers from academia or industry that intend to approach the field.

Workshop: Distributed Embedded Systems *sponsored by Artist*
http://www.artist-embedded.org/artist/-Workshop-Distributed-Embedded-.html *Nov 21st, 2005*
The workshop will have two parts: The first two days are devoted to a limited set of scientific presentations and discussions concerning three major areas: Modular design strategies for distributed embedded systems, predictability and efficiency, design space exploration and application scenarios. The third and fourth day are devoted to Performance Analysis of Embedded Systems.

OSPERT 2005 *sponsored by Artist*
http://www.artist-embedded.org/artist/-OSPERT-2005-Operating-Systems-.html *July 5th, 2005*
Workshop on Operating Systems Platforms for Embedded Real-Time applications. This workshop is intended as a forum for researchers and practitioners of RTOS to discuss the recent advances in RTOS technology and the challenges that lie ahead.

HSCC '05 - Hybrid Systems: Computation and Control *sponsored by Artist*
http://www.artist-embedded.org/artist/-HSCC-05-Hybrid-Systems-Computation-.html
March 9th, 2005
The Eighth International Workshop on Hybrid Systems : Computation and Control (HSCC 2005) attracts researchers from academia and industry interested in modeling, analysis, and implementation of dynamic and reactive systems involving both discrete and continuous behaviors.

6.3 Course Materials available online

The Artist2 NoE participants are massively involved in teaching university courses on topics related to embedded systems. The Artist2 Network of Excellence makes pointers to course materials from its own events and others, available via its website.
Main page for course materials: http://www.artist-embedded.org/artist/-Course-Materials-.html

ARTIST2 Summer School 2008 in Europe
September 8-12, 2008 Autrans (near Grenoble), France
http://www.artist-embedded.org/artist/Programme,1319.html

The Summer School offers a number of foundational tutorials, accompanied by a selection of lectures on exciting emerging technologies and industrial applications - given by leading scientific and/or industrial experts.

ARTIST2 South-American School for Embedded Systems 2008
August 25-29, 2008 Universidade Federal de Santa Catarina, Florianopolis, Brazil
http://www.artist-embedded.org/artist/New-article,1287.html

We believe the school should be the ground for cross-fertilization between Europe and South-America with an expected mutual high added-value. Therefore, the lectures given by European researchers, will be accompanied by talks and a poster session for participants to present and discuss their ongoing work.

Artist2 Summer School in China 2008
http://www.artist-embedded.org/artist/Artist2-Summer-School-in-China,1541.html
July 12-18, 2008 Shanghai, China

ARTIST2 will organize the 3rd edition of a school on Embedded Systems Design in Shanghai. This year, the school is organized in collaboration with the SEI/ECNU and the LIAMA.

Real-Time Kernels for Microcontrollers: Theory and Practice
http://www.artist-embedded.org/artist/Real-Time-Kernels-for,1540.html
June 23-25, 2008 Pisa, Italy

The course on Real-Time Kernels for Microcontrollers aims to introduce the basic concepts of Real-time Systems targeted to Embedded Systems, which are often implemented using microcontrollers. The course will briefly illustrate the theoretical background of real-time scheduling, resource-aware techniques, and wireless communication based upon the IEEE 802.15.4 protocol.

ARTIST2 Graduate Course on Embedded Control Systems
http://www.artist-embedded.org/artist/ARTIST2-Graduate-Course-on,1539.html
May 26-30, 2008 Stockholm, Sweden

The course during 2008 was the 4th in a successful series of course instances, providing an overview and account of state of the art theory and techniques that address the connection and integration of the areas of Control systems and Embedded systems.

LASER Summer School on Software Engineering
http://www.artist-embedded.org/artist/LASER-Summer-School-on-Software.html
September 9-15, 2007 Elba, Italy

The LASER summer school (Laboratory for Applied Software Engineering Research), organized by ETH Zurich, brings together the concepts and practice of software engineering.

FOSAD 2007
http://www.artist-embedded.org/artist/FOSAD-2007.html
September 9-15, 2007 Bertinoro, Italy

The International School on Foundations of Security Analysis and Design (FOSAD) has been one of the foremost events established with the goal of disseminating knowledge in this critical area. The main aim of the FOSAD school is to offer a good spectrum of current research in foundations of security - ranging from programming languages to analysis of protocols, from cryptographic algorithms to access control policies and trust management - that can be of help for graduate students and young researchers from academia or industry that intend to approach the field.

First European-SouthAmerican School for Embedded Systems
http://www.artist-embedded.org/artist/First-European-SouthAmerican,1187.html
August 21-24, 2007 Universidad Argentina de la Empresa (UADE), Buenos Aires – Argentina

The purpose of the school is to foster the well established and dynamic research cooperations in the field of embedded systems between groups in Europe and South America, by allowing south-american students (specially graduate), to meet european researchers.

Artist2 / UNU-IIST School in China - 2007
http://www.artist-embedded.org/artist/Artist2-UNU-IIST-School-in-China,1188.html
August 1-10, 2007 Suzhou (near Shanghai), China

ARTIST2 will organize, in collaboration with UNU-IIST, the 2nd edition of a school on embedded systems design in Suzhou (near Shanghai).

ARTIST2 PhD Course on: Automated Formal Methods for Embedded Systems
http://www.artist-embedded.org/artist/ARTIST2-PhD-Course-on-Automated,1189.html
June 4-12, 2007 DTU - Lyngby, Denmark

Embedded systems engage into an ongoing, hardly foreseeable, interaction with their asynchronously evolving environment. This fact contributes to the intrinsic complexity of their design and validation.

ARTIST2 Graduate Course on Embedded Control Systems
http://www.artist-embedded.org/artist/ARTIST2-Graduate-Course-on,1190.html
May 7-11, 2007 Lund, Sweden

The objective of the course is to provide an overview of the main principles and technologies for supporting the development of embedded control systems.

Real-Time Microcontroller Systems: OSEK Standard and experiments on µcontroller devices
http://www.artist-embedded.org/artist/Real-Time-Microcontroller-Systems.html
March 26-28, 2007 RETIS Laboratory, Scuola Superiore Sant'Anna, Pisa, Italy

Training course on Real-Time Systems for Microcontrollers: OSEK Standard and experiments on microcontroller devices Organised in conjunction with Evidence Srl

Seminar on Quantitative Aspects of Embedded Systems Schloss Dagstuhl 2007
http://www.artist-embedded.org/artist/Seminar-on-Quantitative-Aspects-of.html
March 4 - 9, 2007 Schloss Dagstuhl, Germany

This Dagstuhl seminar will bring together experts in embedded software design and implementation, model-based analysis of quantitative system aspects, and researchers working on extending formal methods with quantitative system aspects.

ARTIST2 - MOTIVES
http://www.artist-embedded.org/artist/ARTIST2-MOTIVES.html
February 19-23, 2007 Trento, Italy

ARTIST2 Winter School 2007 offers foundational tutorials and lectures on exciting emerging technologies and industrial applications - given by leading scientific and industrial experts.

CASTNESS'07 Workshop and School
http://www.artist-embedded.org/artist/CASTNESS-07-Workshop-and-School,1040.html
January 15-17, 2007 Rome, Italy

Computing Architectures and Software Tools for Numerical Embedded Scalable Systems

Real-Time and Control for Embedded Systems
http://www.artist-embedded.org/artist/Real-Time-and-Control-for-Embedded.html
July 10-14, 2006 Pisa, Italy

First European Laboratory on Real-Time and Control for Embedded Systems

ADSD 2006: Advanced Digital Systems Design
http://www.artist-embedded.org/artist/ADSD-2006-Advanced-Digital-Systems.html
September 25-29, 2006 Lausanne, Switzerland

Design course for multimillion-transistor Systems-on-Chip and other state-of-the-art embedded products. The course spans from purely digital-design topics to some compiler-related issues.

LASER Summer School on Software Engineering 2006
http://www.artist-embedded.org/artist/Practical-Programming-Processes.html
September 17 - 23, 2006 Elba, Italy

The 2006 LASER takes an in-depth look at Practical Programming Processes. Many approaches have been proposed in the past decade, from new advances in object technology to Patterns, Aspects, Extreme/Agile/Lean methods, incremental development, process standards (CMMI etc.), Open Source and several others.

Foundations of Security Analysis and Design
http://www.artist-embedded.org/artist/Foundations-of-Security-Analysis.html
September 10-16, 2006 Bertinoro, Italy

FOSAD 2006: 6th International School on Foundations of Security Analysis and Design

Model-Driven Design for Distributed Real-time Embedded Systems (MDD4DRES)
http://www.artist-embedded.org/artist/Model-Driven-Design-for.html
September 4-8, 2006 Brest, France

A goal of this summer school is to provide participants with the information needed to understand and apply MDE approaches to the development of embedded systems. The summer school will also include lectures from experts in academia and industry on topics related to MDE practices and methods, and to emerging MDA technologies.

First European Laboratory on Real-Time and Control for Embedded Systems
http://www.artist-embedded.org/artist/First-European-Laboratory-on-Real,721.html
July 10-14, 2006 Pisa, Italy

Real-Time distributed embedded systems play a crucial role in our society including several application domains such as automotive, telecommunications, robotics, and multimedia systems. These systems generally work under precise timing constraints, to achieve the required level of performance and predictability. Consequently, embedded systems design requires expertise in several disciplines, including control theory, networking, real-time computing, and operating systems.

ARTIST2 Graduate Course on Embedded Control Systems
http://www.artist-embedded.org/artist/ARTIST2-Graduate-Course-on.html
April 3-7, 2006 Prague, Czech Republic

The objective of the Course is to provide an overview of the main principles and technologies for supporting the development of embedded control systems.

ARTIST2 / UNU-IIST Spring School in China 2006
http://www.artist-embedded.org/artist/ARTIST2-UNU-IIST-Spring-School-in.html
April 3-15, 2006 Xi'an, China

The first ARTIST / UNU-IIST Spring School gathered more than 50 participants, of which approximately 40 were students from the top universities in mainland China.

ARTIST2 Summer School 2005
http://www.artist-embedded.org/artist/ARTIST2-Summer-School-2005.html
September 29 - October 2, 2005 Nässlingen, Sweden

The ARTIST2 Summer School was held in conjunction with the 3rd International Conference on Formal Modelling and Analysis of Timed Systems (FORMATS'05). The Summer School offered a number of foundational tutorials accompanied by a selection of exiting new emerging technologies all given by absolute leading scientific experts of the community.

7. International Collaboration

The Artist2 Network of Excellence has always pursued an extremely active policy for organizing and funding International Collaboration events. All of these events were organized and funded by ARTIST.

WESE'08: WS on Embedded Systems Education *October 23rd, 2008*
http://www.artist-embedded.org/artist/-WESE-08-WS-on-Embedded-Systems-.html
As embedded system designs grow more complex and the time to market diminishes, quality embedded systems education becomes more and more important. This fourth workshop on the subject aims to bring researchers, educators, and industrial representatives together to assess needs and share design, research, and experiences in embedded systems education.

Workshop on Foundations and Applications of Component-based Design (WFCD'2008)
http://www.artist-embedded.org/artist/-Components-2008-.html *October 19th, 2008*
The workshop aims to discuss recent results on component-based design with emphasis on design frameworks for real-time systems encompassing heterogeneous composition and models of computation. The focus is not only on fundamental results but also on their implementation in methods and tools and their concrete application in areas such as automotive, avionics, consumer electronics and automation.

ARTIST2 Summer School 2008 in Europe *September 8th, 2008*
http://www.artist-embedded.org/artist/-ARTIST2-Summer-School-2008-.html
The Summer School offers a number of foundational tutorials, accompanied by a selection of lectures on exciting emerging technologies and industrial applications - given by leading scientific and/or industrial experts.

ARTIST2 South-American School for Embedded Systems 2008
http://www.artist-embedded.org/artist/-ARTIST-2-South-American-School-.html
August 25th, 2008
Second edition of the ARTIST South American School.

Artist2 Summer School in China 2008 *July 12th, 2008*
http://www.artist-embedded.org/artist/-Artist2-Summer-School-in-China-.html
ARTIST2 has organized the 3rd edition of a school on Embedded Systems Design in Shanghai. This year, the school was organized in collaboration with the SEI/ECNU and the LIAMA.

ARTIST2 meeting on Integrated Modular Avionics *November 12th, 2007*
http://www.artist-embedded.org/artist/-ARTIST2-meeting-on-Integrated-.html
Integrated Modular Avionics (IMA) has set the principles of standardized components and interfaces of hardware and software in aircraft, applied for the first time in the development of the Airbus A380.

WESE'07: WS on Embedded Systems Education *October 4th, 2007*
http://www.artist-embedded.org/artist/-WESE-07-.html
This third workshop on the subject aims to bring researchers, educators, and industrial representatives together to assess needs and share design, research, and experiences in embedded systems education.

Foundations of Component-based Design *September 30th, 2007*
http://www.artist-embedded.org/artist/-Foundations-of-Component-based-.html
Discuss recent results on component-based design with emphasis on design frameworks for real-time systems encompassing heterogeneous composition and models of computation.

First European-SouthAmerican School for Embedded Systems *August 21st, 2007*
http://www.artist-embedded.org/artist/-First-European-SouthAmerican-.html
The purpose of the school is to foster the well established and dynamic research cooperations in the field of embedded systems between groups in Europe and South America, by allowing south-american students (specially graduate), to meet european researchers.

Artist2 / UNU-IIST School in China – 2007 *August 1st, 2007*
http://www.artist-embedded.org/artist/-Artist2-UNU-IIST-School-in-China-.html
ARTIST2 will organize, in collaboration with UNU-IIST, the 2nd edition of a school on embedded systems design in Suzhou (near Shanghai).

Artist2 - Foundations and Applications of Component-based Design *October 26th, 2006*
http://www.artist-embedded.org/artist/-Foundations-and-Applications-of-.html
The workshop gathered researchers from computer science and electrical engineering to discuss recent results on component-based design with emphasis on design frameworks for real-time systems encompassing heterogeneous composition and models of computation. Especially frameworks for handling non-functional and resource constraints, design under conflicting dependability criteria, trade-offs between average performance and predictability.

WESE'06 - Embedded Systems Education *October 26th, 2006*
http://www.artist-embedded.org/artist/-EmSoft-06-Workshop-on-Embedded-.html
This second workshop on the subject aims to bring researchers, educators, and industrial representatives together to assess needs and share design, research, and experiences in embedded systems education.

ARTIST2 / UNU-IIST Spring School in China 2006 *April 3rd, 2006*
http://www.artist-embedded.org/artist/-ARTIST2-UNU-IIST-China-School-.html
The first ARTIST / UNU-IIST Spring School gathered more than 50 participants, of which approximately 40 were students from the top universities in mainland China.

Artist FP5 : **ACM TECS - Special Issue on Education**
http://www.artist-embedded.org/artist/ACM-Special-Issue-on-Education.html *August 2005*
This special issue of the ACM Transactions in Embedded Computing Systems aims to provide the basis for integrated undergraduate and graduate curricula covering the essential areas of knowledge for tomorrow's embedded systems engineers and researchers.

Artist FP5 : **IST-NSF International Collaboration Days**
http://www.artist-embedded.org/docs/Events/2005/IST-NSF/ *July 7-8, 2005*
- Component-based Engineering for Embedded Systems
- Transatlantic Research Agenda on Future Challenges in Embedded Systems Design

Artist FP5 : **Artist International Collaboration Days 2003 – Education**
http://www.artist-embedded.org/artist/Artist-International-Collaboration,451.html *Oct 11th, 2003*
This was an open meeting to discuss important action lines in the area of Embedded Systems - in which strong synergy between international teams had the greatest benefits. Work over the first year had concentrated on discussion between top researchers in the field, summarized in white papers that were presented and discussed here.

8. Interaction with Industry

Artist2 has a wide array of affiliated industrial and SME partners. Most of these partners have participated in some way in the Artist2 technical meetings and the overall effort. There is strong, high-level industry participation through the various Spreading Excellence events organised by Artist2. Our active involvement in the European Technology Platform ARTEMIS also could have a significant and long-term impact.

In addition, each Artist2 partner has an outstanding track record for interaction with industry. Globally, the Artist2 consortium has a very strong impact on European R&D in embedded systems. This impact is visible via the achievements in STREP and Integrated Projects.

We believe that the strong involvement of four main Artist2 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.

8.1 Interaction with the automotive industry

Specific effort has been dedicated to interacting with the automotive industry. Recall that the automotive industry is one of the two driving sectors for drastic changes to embedded systems design methods, and is certainly *the* sector where changes have been deepest and quickest. This effort was made possible thanks to prior personal strong ties that some key participants (including: Werner Damm (OFFIS), Alberto Ferrari and Alberto Sangiovanni-Vincentelli (PARADES), Martin Törngren (KTH), Rolf Ernst (U. Braunschweig), Sébastien Gérard (CEA)), and affiliates (including: Stefan Kowalewski (RWTH Aachen)) of ARTIST2 had with the Autosar consortium. Albert Benveniste (INRIA) and Werner Damm (OFFIS) jointly organized the *ARTIST2 Workshop* ***Beyond Autosar***[3], held in Innsbruck on March 23-24 2006. The workshop discussed in particular issues related to timing in the Autosar model (the so-called *timing model* of Autosar). More generally, the workshop helped making the academic community aware of the research issues raised by this approach from automotive industry. An elaboration of the results has been presented at EMSOFT 2006 and at a GM Workshop in Bengalore (January 2007).

OFFIS has become a development member of Autosar. This move was proposed to OFFIS by BMW, following in depth technical discussion on the link between the SPEEDS HRC meta-model and the Autosar meta-model regarding timing and safety aspects.

The integrated project SPEEDS has developed a layered meta-model of heterogeneous rich components (HRC) and standardized approaches for the integration of commercial industry standard modelling tools to assemble system-level design models with rich interface specifications by combining models expressed in any authoring tool compliant to the integration standard. A SPEEDS Automotive Day was organized to discuss with the automotive industry how the AUTOSAR methodology can be supported by SPEEDS technologies striving to reconcile the advantage of early system-level analysis with the overall AUTOSAR objective of decoupling function design from its implementation. These results have been presented in several highly visible events, including the DATE 2008 Automotive Day, and a keynote presentation at the Annual Mathworks Automotive Conference 2008 in Stuttgart. More in depth technical discussion on the relation between Speeds HRC model and Autosar were conducted at meetings with BMW, Bosch, and Daimler; see also section on standardization.

Sébastien Gérard (CEA) and Henrik Lönn (Volvo Tech) are organizing a workshop in the context of the ATESST project which aims are inviting key persons working on the context of automotive domain in order to share experience on the usage of standards like MARTE, AADL

[3] http://www.artist-embedded.org/FP6/ARTIST2Events/Events/Innsbruck06/

and Autosar especially in the context of the Architecture Description Language for Automotive, EAST-ADL. They also have setup a new project, ADAMS, (in collaboration with Laurent Rioux from Thales and Julio Medina from Cantabria) dedicated to promote the usage of MARTE in the context of the automotive domain. Let's notice also that this project has to deal also with the aeronautics domain.

MDH has strategic long-term co-operations with seven companies: ABB Corporate Research, ABB Robotics, Arcticus Systems, Bombardier Transportation, CC-Systems, Ericsson, and Volvo Construction Equipment. In recent years, the cooperation has been extended to international subsidiaries and partners of these companies. In addition to this strategic cooperation we have cooperation also with several other Swedish and international companies[4]. The strategy has resulted in substantial industrial support, including a 9.6 MSEK donation from ABB, close to 30 graduate students funded by industry, an industrial lab (currently been set up in cooperation with Ericsson and ABB), a top-talent program for recruitment of international master level students, Adjunct industrial professors from ABB and Volvo, as well as a large number of national and international joint research projects. The cooperation includes the following concrete results:

- further development of the Rubus component model inspired by the Save component model and its implementation in, e.g., Volvo Construction Equipment,
- development of component repository used at CC systems, including components and additional artefacts such as requirements, models and implementations, tests,
- model extraction tools used experimentally at ABB Robotics for modelling real-time properties of legacy systems,
- introduction of model-based approaches for modelling and developing applications in Ericsson,
- work on software decomposition of legacy systems and transformation of development to product-line development at ABB Substation Automation,
- development of a real-time database in cooperation with Mimer,
- in cooperation with core partners building a master program in industrial software engineering focussing on design, architectural analysis, component-based development and dependability of embedded systems.

PARADES has tight links with the ST automotive division and with the Joint Development Group ST-FreeScale and has helped in defining roadmaps for design methodologies, tools and architectures for fault tolerant products. It has a number of interactions with Tier 1 companies including Bosch and Nippon Denso on this very topic. Alberto Sangiovanni Vincentelli is a member of the GM Science and Technology Advisory Board and has fostered joint work with General Motors on distributed embedded system design. In addition, PARADES has contributed to the design of an advanced intelligent component in a tire that consists of a set of sensors, a computing engine, an energy scavenger and wireless communication with Pirelli.

8.2 Interaction with the aeronautics industry

Specific effort has been launched to interacting with the aeronautics industry. This effort was made possible thanks to prior personal strong ties that some key participants (including:

[4] Additional national cooperation includes PhD-students funded by Ardendo, Level21, Prevas, and Scania, as well as joint projects with around 10 SMEs; internationally we cooperate both with giants, such as Nokia, Philips, and Tata, as well as with SMEs, such as Symtavision, Absint, and Rapita Systems.

Werner Damm (OFFIS), Albert Benveniste (INRIA), and Paul Caspi (Verimag)) had with this industry in EU. RTC cluster felt that it was important that the research community around ARTIST2 was made aware of the scientific and technical issues raised by the move to Integrated Modular Avionics (IMA) approach. Recall that the aeronautics industry is one of the two driving sectors for drastic changes to embedded systems design methods, and is certainly *the* sector where changes are most demanding.

Albert Benveniste (INRIA) and Paul Caspi (Verimag), in tight cooperation with John Rushby (SRI, Stanford), have organised an ARTIST2 workshop on IMA, held on November 12-13 2007 in Rome at PARADES location. Speakers include key persons from Airbus, Dassault-Aviation, Israeli Aerospace Industries, Honeywell and WindRiver, plus John Rushby and ARTIST2 participants.

Verimag has recently started a direct collaboration with the European Space Agency ESA which has the objective to adapt results of the OMEGA and the ASSERT project to the needs of the engineers at ESA. A first step consists in an adaptation of the IF tool for UML to UML 2 and to the current version of Rhapsody.

8.3 Interaction with the consumer electronics industry

Thanks to the International Collaboration Days organized within the ARTIST2 project, the ART cluster got in contact with two major companies, Philips and Ericsson, acting in the domain of consumer electronics. After a tight interaction with the engineers responsible for the software development process, a number of industrial needs have been identified, that would make new generation products more robust and flexible.

To cope with a constantly increasing complexity of software applications (already consisting of several million lines of code and hundreds of concurrent activities), a system supporting memory and temporal protection would allow safely mixing real-time and non real-time applications with the benefit of achieving a more scalable platform. Therefore, the work on resource reservation carried out within the ART cluster is of crucial importance to manage the increased complexity of the applications in this domain.

In addition, multimedia systems exhibit a highly dynamic behaviour, since task execution times are often dependent on input data that are difficult to predict. As a consequence, these systems are prone to intermittent overload conditions that could degrade the performance in an unpredictable fashion. Again, the expertise existing in the ART cluster on overload management is of high interest for these companies, since it allows building flexible as well as predictable real-time systems that can react to load changes and perform QoS adaptation in a controlled fashion.

8.4 Interaction with the electronics industry

A new interaction of the ART cluster with Microchip Technology has been started on real-time embedded platforms for monitoring and control. In particular, the expertise existing in the ART cluster on real-time embedded control applications and real-time operating systems is extremely attractive for Microchip, who is interested in pushing the development of real-time embedded applications using 16-bit microcontrollers (as the dsPIC30 and the dsPIC33).

In this context, a big opportunity for the ART cluster is to find an agreement with Microchip to define the characteristics of a small real-time embedded platform for sensory acquisition and motor control that can be used (in conjunction with a wireless card) as a node of a mobile wireless network. This unit would be more powerful and flexible than a mote and could be used to carry out experiments on sensor networks, embedded control, mobile robot teams and distributed control systems.

8.5 Other Cross-sectorial Interaction with Industry

Since April 2008, the INRIA team has started a cooperation with an electronics faculty and a local SME (DeltaDore). This company is a national leader in the domain of home and industrial building equipment. A framework for cooperation has been set up in order to transfer know-how on timed component based architectures. This domain of industry is a promising field for the dissemination of embedded, soft real time component architectures. The challenges of this field lie in the frequent evolutions of deployed architectures. These evolutions call for self configurable and self adaptable components.

In future and many state-of-art projects a convergence of different application domains can be observed for different industrial applications (for example, a multimedia system and safety-critical functions are integrated in a car). In January 2009, the EU STREP project GENESYS (http://www.genesys-platform.eu/) has started with coordination by TU Vienna. The objective of the GENESYS project is to develop a cross-domain reference architecture for embedded systems that can be instantiated for different application domains to meet the requirements and constraints documented in the ARTEMIS strategic research agenda. These requirements are composability, networking, security, robustness, diagnosis, integrated resource management and evolvability. The project will result in a conceptualization of the cross-domain architecture, a specification of cross-domain core services and optional services for the selected application domains, and four exploratory prototypes that will demonstrate and help to evaluate the feasibility of selected central architectural concepts in the different application domains. The analysis of the requirements and the definition of an architectural style with fundamental principles for cross-domain embedded systems have been completed. The next steps will be the definition of the architectural services, the completion of the methodology framework, the implementation of the prototypes and the assessment of the architecture.

The Inria Triskell team is now part of the S3 (Software Services and Systems) European network of excellence. This network started in March 2008 and will end in 2012. The Inria Triskell team is involved in two joint research activities: adaptation and monitoring principles, techniques and methodologies for service based systems, and End to end quality provision and service level agreement conformance. The Triskell team intends to adapt results gained from Artist cooperation on timed components and use these results in the S3 collaborations in the joint research activities mentioned above. The crossbreed between components for embedded systems (Artist2) and service based architectures (S3) will be supported by experiments in the building automation industrial field. Software architectures in this field require real time, reliability, predictability as well as openness and dynamic reconfiguration.

Within the German Competence Cluster SafeTRANS, two SafeTRANS Industrial focussing on V&V methods and on architecture assessments have been organized, with participants from automotive, aerospace, and rail industries. OFFIS is a founding member of SafeTRANS, with Werner Damm being the SafeTRANS Chairman. SafeTRANS has – through its role as a founding cluster of EICOSE - been as well instrumental in deriving research priorities and subprogram formation for the Artemis Joint Undertaking, see below.

OFFIS is also represented through Werner Damm at the Steering Board level of the German Innovation Alliance on Embedded Systems SPES 2020, which is about to be launched in November 2008. This alliance puts together Academic Institutions and Industrial Stakeholders in Embedded Systems development, providing a foundational basis for applications in multiple industrial sectors, including automation, automotive, aerospace, energy, and medical.

Within the Artemis Innovation Cluster on Transportation, EICOSE, the European Expert Group on Transportation has in several meetings identified research priorities for embedded systems from the perspective of the transportation sector, leading to a proposal of the three candidates for subprogrammes (on cost-efficient methods for the development of safety relevant embedded systems, SP1; on Computing Environments for Embedded Systems, SP5, and on

Human Centred Design for Embedded Systems, SP8) for the Joint Undertaking Artemis. All three proposals for subprogrammes were after modification integrated in the Artemins Multi-Annual Strategic Work-Plan. All subprogrammes are cross-sectorial in nature, addressing in particular all transportation sectors. From Artist2, Werner Damm from OFFIS as well as Didier Juvien from CEA are members of the Eicose Steering Board, with Werner Damm serving as EICOSE chairman until May 2008.

8.6 Involvement in ARTEMIS

Several RTC Cluster partners, including CEA, INRIA, OFFIS, PARADES, VERIMAG; and TU Vienna, are actively involved in ARTEMIS, an initiative to form a European technology platform on embedded systems supporting the needs for various industrial and academic embedded application domains, such as the automotive, avionics, but also the real-time requirements of consumer electronics. The interaction with ARTEMIS is expected to influence the work within ARTIST2 positively towards establishing a well-defined conceptual fundament that is useful for academia and industry. Several partners (CEA, INRIA, OFFIS) are involved in EICOSE, the recently established European Institute for COmplex and Safety Critical Embedded Systems Engineering pushed by two French clusters System@tic Paris Région and Aerospace in association with the German cluster SafeTrans. EICOSE has been selected as the ARTEMIS Innovation Cluster on Transportation. VERIMAG and FT R&D contribute within French MINALOGIC cluster to promote the creation of a center of excellence in ARTEMIS encompassing "Nomadic environments" and "Private space" application contexts of the ARTEMIS SRA chart. Contacts have been taken with Nokia and ElectroBit from the Finlandais Symetra Consortium.

The Joint Undertaking Artemis has been created in February 2008, and is about to close the evaluation of proposals submitted to the first call of the Artemis Joint Undertaking. As described above, EICOSE has played an instrumental role in coordinating the research prioirities of both industrial and academic stakeholders in the transportation domain, contributing significantly to the Artemis Multi Annual Strategic Plan in the formation of three out of eight subprograms also forming the basis for the 1st call. Both OFFIS and CEA are members of the EICOSE steering board.

9. Cluster-Level Dissemination and Use of Knowledge

At the cluster level, the Artist2 partners have had significant interaction for dissemination and use of knowledge, both with other clusters and with top teams outside the NoE.

9.1 Real Time Components (RTC) cluster

Interaction with other ARTIST2 Clusters

Since heterogeneity, as well as component-based modelling and analysis naturally involves different aspects of design, then different sub-communities of embedded systems area are interested in this subject, e.g., control, real-time, and hardware. Therefore, RTC topics are a crossing point for several ARTIST2 clusters, in particular RTC, Adaptive Real-Time, Execution Platforms, Control for Embedded Systems, and Verification and Testing. We provide examples of interactions with these clusters.

- **Execution Platforms:** the RTC cluster (Cantabria, EPFL, INRIA, OFFIS, Timisoara, Uppsala, VERIMAG) and the Exeuction Platforms cluster, jointly organized the workshop on "Models of Computation and Communication (MoCC)", organized at ETH Zurich on Nov. 16-17, 2006, by Albert Benveniste, Paul Capsi, and Lothar Thiele (http://www.artist-embedded.org/artist/MoCC-06.html). The objective of the workshop was to survey the different activites in progress concerned with MoCC , gather recognised specialists of the different disciplines in order to attempt to get a panorama of the models used by each discipline, their commonalities and differences, and the several attempts that have already been proposed in order to merge these concepts within some unified view. The problem of compositional analysis of timing and resource properties has been the topic of collaboration during ARTIST2. In 2008, it resulted in a joint publication between ETHZ and Uppsala at EMSORT 2008. **T**here are several new collaborations between the execution platform and the Real-time components cluster. Verimag and ETHZ have collaborated on a translation from the analytic DOL performance evaluation tool to the executable BIP. The newly started Combest project has partners from both clusters (Verimag, EPFL, INRIA, OFFIS, Parades for RTC, and ETHZ and TU Braunschweig for the execution platforms). OFFIS and TU Braunschweig have coorganised a workshop on the certification of safety critical software controlled systems.
- **Control:** Several RTC partners (INRIA, PARADES) are prominent members also in the control community (and some of them are members of the HyCON NoE). Several interactions between the components and control cluster exist. Partners of the control cluster were important contributors to the opening day of the workshop **Beyond Autosar**, which was dedicated to the interaction of distributed embedded software and contro**l.** In year 3, 2 workshops have been organised by participants of different platforms; two of them (at DATE and at CAV) involved partners from the control and component cluster. There are also collaborative projects involving partners from both clusters. In particular, Martin Torngren (KTH) is a core partner in the Control for Embedded systems cluster and collaborates on the "safety critical" platform of the component cluster, in particular through participation in the ATESST project. Short visits have been organised between KTH and CEA in 2008, in order to present work on model based engineering as well as software platforms, and to explore possibilities for joint work.

- **Adaptive Real Time:** TU Vienna is interacting with the group of Eduardo Tovar (ISEP-IPP). Björn Andersson and Rene Cunha from ISEP-IPP participated in the Workshop on Basic Concepts in Mobile Embedded Systems held at TU Vienna, November 2006. Wilfried Elmenreich from TU Vienna was visting ISEP-IPP from May to June, 2007. University of Cantabria has had a fruitful interaction with the group of Luis Almeida (Universidade de Aveiro/IEETA) in the integration of the distributed capabilities provide by the IST-FRESCOR project to the FTT-SE flexible network resource. This continues the effort by Ricardo Marau from Aveiro/IEETA after his visit to Cantabria in 2006. Wilfried Elmenreich submitted his habilitation thesis on “Time-Triggered Transducer Networks” and got awarded the habilitation in 2008.
- **Verification and Analysis:** The very essence of the component platform activities is to integrate component based development with validation. Several cluster partners are also active in the domain of verification and have already good connections to this community. Also several projects, such as the French OpenEMbeDD, the German AVACS, the IP Speeds, the forthcoming COMBEST connect (1) teams working on modelling and model transformation techniques and semantic frameworks and (2) teams working on verification algorithms (3) teams from the execution platform cluster. In 2007, two workshops have been organised jointly with verification and analysis platform partners (see interaction with control cluster). This year, one workshop at least puts a strong aspect on modelling and verification, the workshop on UML and Formal methods. But most workshops focussing on modelling aspects, also include topics on validation.
- **Multi-cluster interactions:** the ARTIST summer schools organised in China and in France involved topics and speakers from several clusters. The topics covered by the summer school in France include Modeling and Validation, Compilers and Timing Analysis, Adaptive Real Time Systems, Control for Embedded Systems, Execution Platforms and MPSoC.

Organization of summer schools

The RTC cluster has been strong drivers in the organization of summer schools

- The Summer School on Model Driven Development for Real-time and Embedded Systems (www.mdd4dres.info) in Sept. 2006 in Brest. This was the third edition of this summer school which focuses on model-driven related issues in the context of real-time and embedded systems development. A new edition of this summer school will be held in Spring 2009 in Autrans.
- An ARTIST Summer School on embedded systems design has been organised in in Shanghai, July 12-18 2008, in collaboration with the SEI/ECNU and the LIAMA. This is the third Artist summer school organised jointly with China. The aim is to promote collaboration between European and Chinese research community on embedded systems and related areas. In 2009, a third edition of this summer school is planned.
- An Artist summerschool has also been organised in Autrans, France in September 2008 (http://www.artist-embedded.org/artist/ARTIST2-Summer-School-2008.html). It is the fourth such summer school organised by Artist in Europe, and it is meant to be exceptional in terms of both breadth of coverage and invited speakers. The topics covered in this year’s school include Modeling and Validation, Compilers and Timing Analysis, Adaptive Real Time Systems, Control for Embedded Systems, Execution Platforms and MPSoC. A balance is seeked between foundational aspects and applications. Speakers include recognized leading researchers and engineers.

Organization of conferences, workshops, summer schools

The RTC cluster has been co-organizing the following conferences and workshops (for more details: see the deliverable on *Spreading Excellence.*

- Sébastien Gérard (CEA) is also co-organizer of a series of workshops on the UML and AADL. The last edition was hold in conjunction with the 13th IEEE International Conference on Engineering of Complex Computer Systems (ICECCS 2008) in *April 2nd, 2008, in Belfast, Northern Ireland (*http://www.artist-embedded.org/artist/Registration,1370.html*)*.
- Sébastien Gérard (CEA) is also co-organizer of a workshop on UML and formal methods: http://www.artist-embedded.org/artist/New-article,1486.html. This workshop will be held in conjunction with the 10th International Conference on Formal Engineering Methods, ICFEM 2008 (October 27th, 2008, Kitakyushu-City, Japan).
- CEA and INRIA are main organizers of the series of Workshop, MODEVVA (www.modeva.org). The objective of this workshop is to offer a forum for researchers and practitioners who are developing new approaches to V&V in the context of MDE. The workshop will discuss V&V of model transformations and code generation; techniques for validating a model or generating test cases from models including simulation, model-checking, and model-based testing; application of MDE to validation, testing, and verification; tools and automation; case studies and experience reports. In 2006, the MoDeVa workshop was been co-located with the MODELS/UML conference in Genova (Italia), and in 2007, in Nashville. In 2008, it has been collocated with ICST'2008, Lillehammer, Norway: CEA and Supélec organized the MoVaH 2008, Workshop on Modeling, Validation and Heterogeneity co-located with ICST 2008, Lillehammer, Norge (wwwdi.supelec.fr/fb/MoVaH08)
- VERIMAG is also a co-initiator and co-organiser of the symposium on Formal Methods for Components and Objects FMCO (http://fmco.liacs.nl/fmco07.html) the aim of which is to bring together researchers and practioners in the areas of software engineering and formal methods to discuss the concepts of reusability and modifiability in component-based and object-oriented software systems The 5th issue has been organised in November 2006 in Amsterdam; In 2007 has taken place a special issue bringing together groups of a set of related EU projects and NoEs; Artist2 is one of those groups. In 2008, the organisation has been taken over by a new group of people.
- Albert Benveniste (INRIA) and Paul Caspi (Verimag), in tight cooperation with John Rushby (SRI, Stanford), and with local support by Alberto Ferrari and Alberto Sangiovanni Vincentelli (who chaired the meeting) at PARADES have organized an ARTIST2 workshop on Integrated Modular Avionics (IMA), held November 12-13 in Rome at PARADES location. Speakers included key persons from Airbus, Dassault-Aviation, Israeli Aerospace Industries, Honeywell and Windriver, plus John Rushby and ARTIST2 participants. See http://www.artist-embedded.org/artist/-ARTIST2-meeting-on-Integrated-.html
- Tom Henzinger, EPFL, and Werner Damm, OFFIS, have organized the Second International Workshop on Foundations of Component-based Design, held at the Embedded System Week in Salzburg on Sep 30, 2007.
- CEA LIST (Christophe Gaston) and Supélec (Frédéric Boulanger) co-organized the MoVaH workshop (http://wwwdi.supelec.fr/fb/MoVaH08/) at the ICST 2008 conference in Lillehammer in April 2008. The topic of this workshop was the modelling and validation of heterogeneous systems.
- MdH organized the workshop COMES'08) on Component Models for Embedded System. This was a workshop with limited number of attendance to foster intensive

interactions, supported by the PROGRESS research centre, MRTC, and ARTIST2. Its aim was to present and discuss the current research and practical results in development of embedded system using component-based development approaches, as well as to discuss and point out the challenges and possible solution directions in applying the component-based approach to achieve predictability of component-based embedded software systems. The workshop was lively, with over 30 participants from PROGRESS and the international research community. http://www.mrtc.mdh.se/progress/COMES/

- PARADES and EPFL organized the workshop: From Embedded Systems to Cyber-Physical Systems: a Review of the State-of-the-Art and Research Needs on Monday, April 21, 2008 in St. Louis, MO, USA. The theme of the workshop was presenting an overarching view of methodologiesand theories for the design of embedded and critical systems as it has emerged in the past five years and discussing the future in terms of the extension of the notion of embedded systems to Cyber-Physical Systems (CPS). In the overview of the present status of the discipline, the workshop addressed heterogeneous system composition, design methods based on abstraction and refinement, interface theories, mapping of abstract entities to implementation platforms and industrial applications. The presentations featured industry representatives who gave their perspective of what are the gaping holes in the state of the art in their business segment and how to bridge academic accomplishments with industrial practice. The discussion about the extension of the theories and methodologies to the new generation of CPS reviewed the necessary steps and a possible roadmap for research. The discussion included public research organizations. European Community representatives, Werner Damm as Autosar and Artemis representative and Philippe Reynaert, DG INFSO Embedded Systems, provided the state-of-the-art and the research initiatives on embedded systems in the EU.
- OFFIS and TU Braunschweig have organised the workshop “SafeCert 2008 – Certification of Safety-Critical Software Controlled Systems” (http://safecert08.offis.de) as a satellite event of ETAPS 2008. The major question addressed in the workshop was how to embed formal methods and tools in a seamless design process which covers several development phases and which includes an efficient construction of a safety case for the product.
- Alain Girault and Eric Rutten (INRIA) organized the Model-driven High-level Programming of Embedded Systems, SLA++P’08 (a satellite event of the European Joint Conference on Theory and Practice of Software, ETAPS 2008), in Budapest, Hungary, March 2008. The keynote speaker was Grégoire Hamon from Mathworks. See http://www.artist-embedded.org/artist/SLA-P-2008,1231.html
- Thierry Jéron (INRIA) co-organized the MOVEP’08 summer school on modeling and verifying parallel processes, in Orléans, France, June 2008. See http://www.univ-orleans.fr/movep2008/
- In the context of the ATESST project, CEA and KTH organised two workshops. The “EAST-ADL, AADL, MARTE, Autosar harmonization workshop” which provided useful information exchange between the project and the respective standardization initiatives. It was agreed to maintain contacts, and to organize follow up meetings. Identified topics of common interest include Timing, Error modeling and Methodology. The second one is on the “Model based development of automotive embedded systems - The EAST-ADL approach” (march 3, 2008 - www.md.kth.se/RTC/atesst-open-workshop_v1.1.pdf).
- CEA co-organises the workshop on UML and AADL is held in conjunction with the 13th IEEE International Conference on Engineering Complex Computer Systems, ICECCS 2008, this workshop gathered researchers and practitioners interested in all aspects of

the representation, analysis, and implementation of DRE system behaviour and/or architecture models. http://www.artist-embedded.org/artist/-UML-AADL-2008-.html

- CEA and Verimag coorganised (with external partners) the 1st International Workshop on Model Based Architecting and Construction of Embedded Systems held in conjunction with MODELS 2008 as a follow-up workshop of the SVERTS and MARTE workshops organised in previous years, the objective of this workshop is to bring together researchers and practitioners interested in model-based software engineering for real-time embedded systems. We were seeking contributions relating to this subject at different levels, from modelling languages and semantics to concrete application experiments, from model analysis techniques to model-based implementation and deployment. Given the criticality of the application domain, a particular focus is on model-based approaches yielding efficient and provably correct designs. http://www.artist-embedded.org/artist/ACES-MB-08.html
- CEA co-organises (with external partners) the 1st IEEE International workshop UML and Formal Methods as a satellite of ICFEM 2008 in Japan. For more than a decade now, the two communities of UML and formal methods have been working together to produce a simultaneously practical (via UML) and rigorous (via formal methods) approach to software engineering. The fact that the UML semantics is too informal has led many researchers to formalize it with different existing formal languages. The main objective of this workshop is to encourage new initiatives of building bridges between informal, semi-formal and formal notations.

  http://www.artist-embedded.org/artist/UML-FM-08.html
- MdH and Uppsala (via ARTES and SNART) have driven the organization of the Swedish Embedded Systems Meeting in Stockholm, March 5, 2008, where results of Swedish research programs on Embedded systems are presented, and further developments are discussed. http://www.snart.org/conference/2008/ses/

9.2 Adaptive Real Time (ART) cluster

Interaction with the control community

The ART cluster had several interactions with the control community and in particular with the cluster on Control for Embedded Systems. Since the first year, the two cluster leaders, Giorgio Buttazzo (ART) and Karl-Erik Arzen (Control) organized a number of meetings and workshops to exchange ideas and propose more concrete actions to make progress in this area.

A joint work involving people from Pavia, Pisa and Lund has been carried out to integrate feedback control schemes into the Shark operating system (used as a shared platform) and to investigate the effects of different scheduling policies on delays and jitter in control loops.

Another strong collaboration has been established with the hybrid systems community. As a result of this connection, Giorgio Buttazzo has been invited as a co-Program Chair to organize the International Conference on Hybrid Systems: Computation and Control (HSCC 2007).

A joint work involving people from UPC (affiliated to TUKL) and Lund has been carried out to investigate feedback scheduling techniques. A PhD student from UPC spent 5 months in Lund working on the project.

In the last year, Pisa, Lund and TUKL started a collaboration to achieve adaptive resource reservations in multi-core systems. Pisa contributed to identify the most appropriate scheduling algorithms, Lund contributed on feedback control schemes, and TUKL on defining the programming interface.

Interaction with the cluster on compilers and timing analysis

A collaboration has been started with the cluster on compilers and timing analysis to investigate the problem of enhancing the predictability of real-time systems by reducing the variability of task execution times. In fact, internal kernel 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.

To address these issues, a new research was initiated that looks at predictability and efficiency in a synergistic manner and that involves all levels of abstraction and implementation in embedded-system design.

Thanks to the ARTIST2 network of excellence, the ART cluster got in contact with the cluster on Compilers and Timing Analysis. The two clusters started working together to develop a new approach consisting of a combination of several methods, including (a) design-space exploration on the hardware architecture level to identify good designs offering combinations of strong performance with good predictability, (b) appropriate kernel mechanisms for task and resource management that are predictable and analyzable, and (c) a synergistic development of models, design methods and matching analysis tools that extract precise system-behaviour properties.

Interaction with the language community

The ART cluster participated in the development of Ada, (Ada 2005), Java (RTSJ) and POSIX (for use with C and C++). This participation has included membership of the associated standardisation bodies that linked the work within the cluster with international efforts across such languages.

Interaction with the real-time components community

A collaboration between the clusters on components and adaptive real-time has been carried out along the ARTIST2 project. The main goal is to provide support for dealing QoS aspects in component-based systems. This technology is a relevant approach to complex system development and to allow a smooth integration of software from different vendors. QoS management is an adequate mean to provide a predictable quality to end-users. The collaboration between those clusters has brought competencies in component-based design for hard and adaptive real-time systems, to produce advances that would be difficult to achieve without all three.

This cooperation has facilitated the development of a number of technical achievents along four research lines: a) specification of QoS properties using UMNL profiles and aspect-based approachs, b) generation of analyzable models from the UML models, c) composition of QoS-aware components and adaptability, and d) QoS support in run-time components frameworks. The participants in this activity have actively participated in the development of a number of OMG standards

Dissemination

The ART cluster has been quite active in disseminating the research results achieved in the context of the ARTIST2 network of excellence, as an overall strategy for reaching other research/academic/industrial communities with related interests.

The operating system platform developed in the context of the Joint Programme of Integration Activities (JPIA) has been extensively used in summer schools and graduate courses to teach how to develop embedded applications with real-time and performance requirements.

In additions, several scientific papers have been published and a number of workshops, conferences, and invited talks have been organized by the ART cluster to spread the acquired knowledge in the scientific community. The conferences and workshops in which the ART cluster has been involved include:

- OSPERT 2008: Workshop on Operating Systems Platforms for Embedded Real-Time applications, Prague, Czech Republic, July 1, 2008.
- ETFA 2007: IEEE International Conference on Emerging Technologies and Factory Automation, Patras, Greece, September 25-28, 2007.
- RTSS 2007: IEEE Real-Time Systems Symposium, Tucson, Arizona, USA, December 3-6, 2007.
- ECRTS 2008: Euromicro Conference on Real-Time Systems, Prague, Czech Republic, July 2-4, 2008.
- RTAS 2008: IEEE Real-Time and Embedded Technology and Applications Symposium, St. Louis, MO, United States, April 22-24, 2008.
- HSCC 2007: ACM International Conference on Hybrid Systems: Computation and Control, Pisa, Italy, April 3-5, 2007.
- RTCSA 2008: IEEE International Conference on Embedded and Real-Time Computing Systems and Applications, Kaohsiung, Taiwan, August 25-37, 2008.
- IFAC 2008 World Congress, Seoul, Korea July 6-11, 2008.
- IECON 2007: Annual Conference of the IEEE Industrial Electronics Society, Taipei, Taiwan, November 5-8, 2007.
- WFCS 2008: IEEE International Workshop on Factory Communication Systems, Dresden, Germany, May 20-23, 2008.
- RTN 2008: International Workshop on Real Time Networks, Prague, Czech Republic, July 1, 2008.
- WPDRTS 2008: International Workshop on Parallel and Distributed Real-Time Systems (In conjunction with IPDPS), Miami, Florida, USA, April 14, 2008.
- Ada Europe 2008: International Conference on Reliable Software Technologies, Venice, Italy, June 16-20, 2008.
- ISORC 2008: IEEE International Symposium on Object and component-oriented Real-time distributed Computing, Orlando, Florida, USA, May 5-7, 2008.
- WTR 2007: Brazilian Workshop on Real-Time Systems, Belem, Brazil, May 28th, 2007.
- RTNS 2007: Int. Conf on Real-Time and Networked Embedded Systems, Nancy, France, March 29-30, 2007.
- SAE 2008 World Congress, Detroit, Michigan, USA, April 14-17, 2008.

9.3 Compilers and Timing Analysis (CTA) cluster

Cluster members performed teaching activities (e.g. Peter Marwedel from Dortmund and Rainer Leupers from Aachen at ALARI, Lugano and EPFL, Lausanne) in cooperation with

other ARTIST2 members (e.g. Luca Benini/Bologna and Lothar Thiele/Zürich). Further links of ARTIST2 members existed to the SHAPES project and to the HiPEAC Network of Excellence.

TU Dortmund

The interaction with the local technology transfer centre ICD (see http://www.icd.de/index_eng.html) is key for interacting with industry. ICD is headed by Peter Marwedel. ICD is used for transferring research results to industry. The group promoted education in embedded systems through a published text book and through courses at EPFL, at ALARI and at spring or summer schools in New Zealand, China, Brazil, Portugal, Korea, and Germany. TU Dortmund organizes the SCOPES series of workshops on compilation for embedded systems.

RWTH Aachen

A close cooperation existed with the ARTIST2 Execution Platforms cluster, in particular between Dortmund, Aachen, and Bologna University. RWTH Aachen participated in the HiPEAC network of excellence and started new cooperations related to code optimization, e.g. with Edinburgh University. Furthermore, Aachen maintained tight industry cooperations, e.g. with CoWare, ACE, and Infineon. Since Oct 2006, RWTH Aachen is running the UMIC research cluster (http://www.umic.rwth-aachen.de) of the German excellence initiative.

ACE

ACE worked closely with ST and with Philips having both a commercial relationship with them as well as being co-members of EU project consortia – in one case along with Verimag. ACE has been working closely with Aachen in this domain for some time. One of the results of this cooperation has been the integration of compiler technology in a start-up company that span out of the university. Cooperation with Imperial College and Edinburgh has also started.

AbsInt

Within the EmBounded Project (IST-510255), AbsInt was also involved in the development of the Hume compiler. Hume is a domain-specific high-level programming language for real-time embedded systems.

TU Berlin

TU Berlin is generally involved in methods and tools for software engineering for embedded systems. TU Berlin has cooperated with Edinburgh University (Björn Franke) concerning the optimization of compilers based on machine learning techniques. Furthermore, TU Berlin has done research on the verification of embedded operating systems, also by cooperating with the Fraunhofer institute FIRST. Finally, TU Berlin visited and was visitied by other cluster members, e.g. ACE, RWTH Aachen and TU Vienna.

IMEC

IMEC is integrated in European research networks, including HiPEAC. Moreover, IMEC is the central partner of a Marie Curie Host Fellowship project that involves more than 10 universities across Europe. IMEC also has many industry co-operations including most large European multi-media and communication systems oriented companies.

Mälardalen

The WCET analysis group maintains close contacts with several industrial partners, and has conducted a number of case studies using their production codes. The group also interacts heavily with the Component-based Software Engineering community through the national centre PROGRESS for research on component-based software design for embedded systems.

Saarland University

Timing-Analysis activities in the cluster interacted closely with the Execution-Platform cluster in the area of increasing the timing-predictability of real-time systems. Airbus and Bosch participated in the Predator FP7 proposal aimed at reconciling performance with predictability. The PREDATOR project started in 2008. Saarland University worked on modularizations of the Sagiv/Reps/Wilhelm shape analysis together with Mooly Sagiv, Tel Aviv University, and Arnd Poetzsch-Heffter, University of Kaiserslautern.

Tidorum, York

Tidorum (and partially York through Rapita Systems) were engaged in a project for the European Space Agency to study the timing and verification aspects of cache memories in space systems. The PEAL project ended in February 2007. An extension was started in late 2007. The main partner from the aerospace domain was Thales Alenia Space, France.

TU Vienna

TU Vienna worked on measurement-based timing analysis together with TU Munich. It also maintained a close interaction with the Lawrence Livermore National Laboratory, CA, USA, in optimizing high-level abstractions.

See ARTIST2 Y1, Y2, Y3, and Y4 deliverables on activities for a more exhaustive description of the state of the integration.

9.4 Execution Platforms (EP) cluster

ETH Zurich has been organizing and participating in the CASTENESS Workshop, see www.casteness.org. The workhop put together the expertise of various EU projects such as ARTIST2, SHAPES, AETHER. In addition, ETH Zurich has been given a tutorial on issues that have been investigated in the ARTIST2 context: Analytic Performance Estimation, Mapping Algorithms to Architectures, Scalable SW Construction. The workshop has been sponsored by ARTIST2 and took place 15.-18th of January 2008.

ETH Zurich has given part of a summer school/advanced course on ADVANCED DIGITAL SYSTEMS DESIGN on 10-14 September 2007, Lausanne, Switzerland. The participants are from industry and university. This way, results from the integrated view of embedded system design will be brought to a much larger community.

Lothar Thiele from ETHZ has been given a tutorial at EMSOFT on Sept. 30 2007, the major conference in the area of embedded software. It covered methods for performance analysis of distributed embedded systems and presented outcomes of the ARTIST2 project.

As a follow-up of the Models of Computation and Communication (MoCC) at the ETH last year, the TU/e organized the MoCC2008 workhop. It took place in Eindhoven on July 3th and 4th. It brought together scientists from various areas, i.e. formal methods, hardware design and software architecture, http://www.artist-embedded.org/artist/MoCC-2008.html.

The ESI (TU/e) participated in the Quasimodo workshop (IST framework programme 7) and brought in a industrial case study to combine different quantitative analysis techniques. The workshop was held in Aachen on June 2nd and 3th.

TU Eindhoven and TU Braunschweig are guest editors of the ACM TECS special issue on Model-driven Embedded System Design (http://acmtecs.acm.org/mesd.htm).

Linköping has given an invited talk at the DATE 2008 Conference, as part of the special day on Dependable Embedded Systems. With this occasion several results obtained in the ARTIST context have been made accessible to an international audience. They are related, in particular, to fault tolerance aspects of distributed real-time systems like those used in automotive applications.

Linköping has organised the 6th IEEE Workshop on Embedded Systems for Real-Time Multimedia, as part of the ARTIST sponsored Embedded Systems Week 2008.

UNIBO has been very active in the Multi-core Systems-on-Chip community and in the computer architecture community which is now aggressively targeting multi-core systems. UNIBO has become active member of the HIPEAC2 network of excellence and participated to several events in this area. Prof. Benini has been an HIPEAC instructor at the ACACES summer school, in L'Acquila. Members of the UNIBO team have participated to the main HIPEAC events in 2008.

DTU has been organizing the DaNES Mini-Case Workshop on industrial case-studies at DTU on May 22-23, 2008.

TU Braunschweig has been organizing the Embedded Software Track at the major European conference on design automation DATE (Design Automation and Test in Europe) that took place March 10-14, 2007. The track was devoted to modelling, analysis, design and deployment of embedded software, including formal methods, tools, methodologies and development environments. Thereby, the emphasis was on embedded software platforms, software integration and portability issues.

9.5 Control for Embedded Systems (Control) cluster

Similar to previous years the main interaction ¨within Artist2 has been with the ART cluster and the RT-Components cluster. The interaction with the ART cluster has been performed through joint research work, and through joint proposals and projects.

Outside Artist2 the cluster has interacted with a number of other communities. Some examples are given below:

- The partners of the cluster have interacted with the partners in HYCON through joint participation.
- The partners of cluster have interacted with the partners in SOCRADES IP projects through joint participation.
- The partners of the cluster have interacted with the partners in numerous STREP projects. These include ATESST, ACTORS, DYSCAS, CHAT, AEOLUS, and FRESCOR.
- The partners of the cluster have interacted with the respective national research communities.
- The cluster has organized or co-organized a number of workshops and events, both with a research focus and with a dissemination focus. These includes:

- The Fourth Graduate School on Embedded Control Systems, Stockholm, May 2007.
- Zdenek Hanzalek was the General Chair for the 20th Euromicro Conference on Real-Time Systems (ECRTS 08) held in Prague, July 2-4, 2008.
- A one-day workshop on "Embedded Control Systems: From Design to Implementation" was held in association with the IFAC World Congress, Seoul, Korea, 6 July, 2008
- The cluster was among the presenters at the workshop "Complex Embedded and Networked Control Systems" organized by the HYCON community and held in association with the IFAC World Congress, Seoul, Korea, 5-6 July, 2008
- The cluster co-organized a workshop on "DataFlow Modeling for Embedded Systems" together with the ART cluster and the ACTORS project that was held in Pisa, 5 May, 2008
- An invited session on networked embedded control for the CDC 2008 conference was organized together with Albert Benveniste from the RT-Components cluster. The session was accepted and will take place in Cancun in parallel with the final Artist2 review in December 2008.
- Karl Henrik Johansson and Karl-Erik Årzén co-organized the EU-US'08 workshop held in Stockholm, 16 June 2008. The topic of the workshop was Networked Information and Control Systems.
- A workshop on "Model-based development of automotive embedded systems – The EAST-ADL approach" was co-organized with ATESST, Brussels, 3 March, 2008
- A workshop on harmonization of modeling languages was co-organized with ATESST in Paris, 25 October, 2007. The meeting gathered representatives from the EAST-ADL, AADL, and MARTE communities.
- A KTH/Industry seminar was organized on September 3 to mark the kick-off for the KTH Embedded Systems centre (ICES). Presentations were given among others by representatives from ABB, Ericsson, and Scania and by Edgar Brinksma from the Dutch Embedded Systems Institute.

- The partners of the cluster has given several keynote addresses, invited sessions, and invited lectures, see the respective activity reports.

9.6 Testing and Verification (TV) cluster

At the *scientific level* model checking technology forms the very basis for automatic verification with numerous applications. Its recognition in Computer Science as a core technology is clearly witnessed by the giving the Turing Award 2007 jointly to Edmund Clark, Allan Emerson and Joseph Sifakis for their original and continued research on model checking.

In the period of ARTIST2 model-checking has been successfully applied to the automatic generation of test suites (with guaranteed coverage), and is also increasingly applied successfully within and by other communities including hardware/software co-design, control theory, discrete event systems, fault-tolerance, planning and scheduling and performance evaluation.

Members of the cluster has published and given invited talks at main conferences and in journals of these neighbouring communities.

Similarly leading research groups within AI are finding applications of existing search heuristics from planning to the improved model-checking (e.g. Friburg University, Germany within the AVACS project and Trento University, Italy).

At the *organization* level, members of the cluster have been active in the European ARTEMIS initiative; in particular ESI is a member of ARTEMIS, and other partners of the cluster have been active in promoting ARTEMIS at national levels (e.g. Aalborg together with IMM/DTU have been initiators of the Danish D-ARTEMIS consortium).

10. Vision Beyond the Artist2 NoE

10.1 Real Time Components (RTC) cluster

Embedded real-time components are used in more and more application domains, which go clearly beyond avionics and automotive. Building automation is one of these domains, with a predictable growth in the next ten years. This growth is motivated by many factors. Two factors are of particular importance: energy saving and health care for the elderly. The increasing concerns about energy use and pollution has made energy conservation and use of renewable sources of energy a primary goal of the European Community. In this context, we believe that the overall activity of ARTIST 2 is essential since embedded control and monitoring will be a pillar to achieve possibly overambitious goals to have buildings that consume zero net energy by 2012.

The lines of work of the different RTC activities will be continued in many forms.

- As already mentioned in the previous section, the need for a good and effective connection between UML-RT related standardisation bodies and the active academic community still remain. This is the role currently played by the activity *Development of UML for Real-Time Embedded Systems.* This role should remain fulfilled even after the end of ARTIST2 and the effort to disseminate around the MARTE specification will be more and more effective in the incoming year. Both the ADAMS project and the Artist Design NoE will two of the most important support for achieving this purpose.
- The aim of the component platform activity is to show the feasibility of and possibly to improve the design approaches for component based heterogeneous systems in the cluster by providing tool support for it. To this aim, we have started to build a set of platforms or tool suites supporting such model-based design approaches based on user level modelling notations, notations supported in commercial tool suites, and new modelling paradigms being developed in this cluster. At a longer term, the today isolated tool suites showing partial solutions should be usable consistently also in a combined fashion due to the existence of semantically well-founded component frameworks flexible enough to represent and meaningfully combine models from different user tools and possibly different abstraction levels or view points and that can be exploited by back-end tools (analysis and code generation). Generally speaking, most of the activities of the platform activities will be continued in the modelling and validation cluster of the Artist Design NoE in some form, even if the platform will not be identified as such in the new NoE. T
- The challenge of putting model-based design of embedded systems on a firm scientific basis has met with further problems that should be addressed. Important problems include to bridge the dichotomy between operational and transformational modeling approaches. Operational means automata-based: these approaches work on a component level, and have been successful in model checking, protocol verification, and code generation. Transformational means stream-based: these approaches work on the system level, and have been successful in performance analysis. While operational approaches have difficulties to scale to systems, transformational approaches suffer a loss of precision. Further important topics include resource modeling, to permit the exploration of trade-offs between multiple dimensions, such as functionality, reliability, performance, and resource consumption. To overcome the problem that current models for modeling quantitative properties of systems systems (Markov processes; timed automata; hybrid automata) tend to be brittle and overly sensitive towards arbitrarily small numeric perturburances, we plan to develop robust models for stochastic, timed, and hybrid systems. On the tool side, we will continue to

study the integration of BIP and Metropolis as two frameworks that will allow the composition and the analysis of heterogeneous parts. The sequence of meetings in different industrial sectors will continue by considering the building, avionics, consumer electronics and automotive domains (a meeting is planned for November 12th and 13th, 2008 at PARADES to discuss these application domains and see how Artist Partners can contribute to solidify a design methogology and supporting tools for industry. A major focus of future work will be to establish well defined bridges between innovation potentials for component based design and industry standards, notably Autosar.

The above and other challenges will be addressed in current and to-be-initiated collaboration projects between Artist2 partners and others. Examples include ArtistDesign, and the omgoing projects SPEEDS, COMBEST, and GENESYS.

10.2 Adaptive Real Time (ART) cluster

The major result of the ART cluster has been to build a significant amount of knowledge on problems, methodologies, techniques, and tools for embedded systems with highly dynamic behavior. Such a knowledge is now available to be disseminated in the industry and in the academia to educate next generation engineers.

The vision beyond the ARTIST2 NoE is to organize such a huge amount of knowledge and make it available in different forms to help the development of embedded systems that are more robust, more efficient, more flexible, and more predictable than what is possible today.

Part of this cluster will continue to work within the ArtistDesign NoE, both in the OS and Networks domain, and in the transversal Design for Adaptivity activity. A lot of the work will also continue in the different existing and new IP and STREP projects that the core partners are members, e.g., FRESCOR, ACTORS, PREDATORS and INTERESTED, as well as in different national projects.

There are strong indications that adaptive real-time techniques will continue to be important for the embedded systems community. Scheduling and resource management must allow a higher flexibility to handle future applications, which are going to be more dynamic in terms of resource requirements.

The current industrial trend of developing multi-core platforms is introducing a higher degree of complexity that is pushing the research community towards new approaches and methodologies. In fact, the traditional programming model used so far in uniprocessor platforms is quite inadequate for systems consisting of multiple cores and needs to be completely revisited.

10.3 Compilers and Timing Analysis (CTA) cluster

Fortunately, the successful cooperation within the ARTIST2 NoE will be continued in the ArtistDesign NoE. For ArtistDesign, the scope has been extended. There is now the focus on multiprocessor systems, due to the needs of the industry in this area. This includes the integration of the single-task-on-uniprocessor methods. Also, this includes the integration of tools into design flows considering distributed and communication-centric systems. This new direction is considered by both subclusters, as multiprocessor systems generate new problems for compilation and for timing analysis. The subclusters are also cooperating on this issue. A first workshop has been held. Also, code generation beyond compilers receives more attention. Work on compiler platforms, timing analysis, and resource aware design continues.

New projects extending the cooperation between the partners include the following:

- AbsInt, Mälardalen, York spinoff Rapita Systems, and TU Vienna are partners in the FP7 STREP ALL-TIMES (ICT-215068, duration Dec. 2007 – Feb. 2009). Their work on tool integration and analysis of C code will be continued within ALL-TIMES.
- AbsInt, Saarland University, and TU Dortmund are partners in the FP7 project PREDATOR (Ref. 216008, duration Feb. 2008 – Jan. 2011) aiming at reconciling predictability and efficiency. Joint work on the WCET-aware compiler will continue in that framework.
- IMEC, TU Eindhoven and ICD (a spin-off of TU Dortmund) are partners in the FP7 Mnemee (IST-216224, duration Jan. 2008-Dec. 2010) project.

Timing Analysis: code-level timing analysis is now in a state where it is being applied in industry for the analysis of time- and safety-critical systems. Still, many challenges remain. Code-level analysis tools and techniques must be integrated into tool chains and development processes, interacting with system-level tools. The usability can be improved, increasing the level of automation. Early, approximate estimates of timing properties are needed for the dimensioning of systems. As computer architecture develops, the timing models become more complex and the subsequent analysis becomes harder. In particular, the introduction of explicitly parallel multicore and MPSoC architectures is very problematic from a timing analysis point of view since shared resources, like buses and memories, easily can cause very unpredictable timing behaviour.

What is needed is a new design discipline, Design for Predictability, that deals with resource handling in general and cuts across several disciplines such as timing analysis (as well as other resource analyses), computer architecture, software design, system software, compilers, and HW/SW codesign/synthesis..

The ultimate vision is a fully integrated development process with resource needs and safely and precisely determined resource consumption communicated between components and layers through resource interfaces. The PREDATOR project addresses these issues.

Compilers: multiprocessor systems challenge compiler technology. Significant progress is required to support the forthcoming parallel architectures. Significant investments into compiler technology are needed to meet the continuing trend towards more performance hungry applications. The lack of adequate compiler technology could potentially inhibit new applications in a large variety of domains. In many cases, required technologies can use available compilers as backends. Hence, fortunately, there is a decreasing need to integrate all new optimizations into a single monolithic compiler.

The generation of embedded systems from specifications in standard von-Neumann languages comes with a number of problems, like potential deadlocks, priority inversion etc. Therefore, new models of computation are being tried out. New code generation techniques are needed for such new models.

Research along the lines of the ARTIST2 project continues to be needed. New optimization engines are required. Resource aware design will be an important topic in the future. This includes memory-architecture aware compilation in particular. Due to the obvious impact of the memory-wall problem, embedded systems will become memory-speed limited and all techniques easing the problem can be expected to find a major attention. Also, optimization engines for SIMD architectures must be improved. Energy efficiency of embedded systems will become even more important than it was in the past. The required high levels of optimization will need to be supported by advanced code analysis techniques. Code correctness is urgently needed. Hence, techniques for code verification are also needed.

Synergies between compilers and timing analysis: Linking compilers and timing analysis can be expected to be an area of further research. Improved availability of timing information in compilers is certainly overdue. It can be expected that this will be recognized outside this

consortium as well and that timing issues will be given more attention. It is unknown, how quickly this will be taken up by the industry.

10.4 Execution Platforms (EP) cluster

As a result of the activities of the execution platforms cluster, a number of powerfull analyses, design and exploration methods are available today. An important direction for future activities within ArtistDesign is to make a step towards large-scale industrial applications by applying these methods to industrial applications in different domains, analyse their strengths and weaknesses, categorize them, integrate them and fit them in industrial design trajectories.

In ArtistDesign we will continue our work on analysis and optimisation of distributed embedded systems, fault tolerance, and energy efficiency. The focus will be on systems with a dynamic nature where we do not assume that a certain worst case constellation is known at design time, but certain decisions regarding resource allocation have to be taken dynamically, at run-time, by still satisfying certain safety and QoS constraints.

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

- The integration of different levels of Quality-of-Service will be investigated in the AIS Project (http://www.edacentrum.de/ais/)
- In particular, mixing real-time and non-real time processing in a multi-core system is the focus of the Compose project (http://www.ida.ing.tubs.de/en/research/projects/compose/)
- In the scope of the Combest project (http://www.combest.eu), TUBS and ETHZ build on the foundations of hierarchical event models and scheduling.
- Organic computing has become a Schwerpunktprogramm of the DFG (http://www.organic-computing.de/spp)
- Intertask communication and synchronization in multi-core systems.

10.5 Control for Embedded Systems (Control) cluster

The general vision for the research work that is coordinated within the cluster is summarized in the following two statements:

Development of methods, tools and theory that allow faster and more efficient development of networked embedded control systems that are safer, more flexible, more predictable, have higher degree of resource utilization, and better performance than what is possible today

and

Advance the state of the art in applying control methods for providing flexibility and robustness and manage uncertainty in embedded computing and communication systems.

This cluster as a whole will be discontinued after the end of Artist2. However, in spite of this, a lot of the work will still be continued in the ArtistDesign network, both in the OS and Networks cluster and in the transversal Design for Adaptivity activity. A lot of the work will also continue in the different existing and new IP and STREP projects that the core partners are members, e.g., SOCRADES, FRESCOR, DYSCAS, ACTORS, and ATESST, as well as in different national projects.

There are strong indications that control implementation techniques will continue to be important for the embedded systems community. Control is and will without doubt continue to

be one of the largest application areas for embedded system, in particular for ubiquitous networked embedded systems. The current multi-core trend that both makes traditional static implementation techniques more difficult and generates new requirements on programming models and implementation techniques is one sign of this. Another sign is the focus on small ubiquitous networked devices in the form of, e.g., sensor networks, where there are severely limited computing resources, but still a desire to perform as much of the computations (incl. control computations) locally in order to save communication bandwidth and battery power.

The use of feedback to provide performance and robustness in networked embedded computer systems becomes more natural, the more complex and hard to statically analyze the systems are. Since increased complexity and an ever increasing amount of software is one of the most prominent trends in embedded systems today we are convinced that dynamic feedback-based resource management will be increasingly important for the future.

10.6 Testing and Verification (TV) cluster

As clearly observed by the many industrial contacts of the two national embedded systems centers, ESI (The Netherlands) and CISS (Denmark), testing is *by far* the most used and important validation technique applied by industry today. It is estimated that some *30-70% of the total development cost* for embedded systems is spent on testing at various stages. It is also a general observation that current testing practice is very ad-hoc often with manual construction and even execution of test-scripts. There is clearly a gap between current industrial practice and existing academic state-of-the art technology. It is important that continued effort is made towards bridging this gap though collaborative projects attempting to make industry take-up existing state-of-the-art testing and verification techniques.

To focus on aspects such as performance, timeliness, and efficient resource-usage, the testing and verification techniques should be based on models with *quantitative information.* To provide a coherent model-based testing and verification methodology with a well-integrated chain of tools applied in industrial practice is a long-term vision beyond the ARTIST2 NoE. In addition to modelling, this will require a strong focus on analytical techniques that address the combination of non-determinism, real-time and stochastic information. Here we foresee the need for techniques that will combine Abstract Interpretation and Model Checker. Also support for high abstraction levels must be provided to overcome the inherent complexity of modern embedded systems. Finally, (semi-)automatic generation of code that preserves the relevant design properties will be essential to ensure industrial impact. These challenges on quantitative modelling and verification will be key activities within the ARTIST Design NoE.

The momentum and willingness of the partners of the cluster to continue working together is very strong. This is witnessed by the two newly started EU FP7 STREP projects *Quasimodo* and *Multiform*. Here *Quasimodo* is targeted towards quantitative modelling formalisms and tool plug-ins for the use in specification, analysis, implementation of embedded systems, and *Multiform* focuses on tool-integration of academic tools (e.g. PHAVer and UPPAAL) and commercial tools (e.g. Simulink). Other partners of the cluster are active in the SPEEDS project where formats for tool exchange are becoming available.

The partners of the cluster also intend to play an active role in the forth-coming Joint Technology Initiative ARTEMIS' research priority on Design Methods and Tools. Here ESI already play a leading role.