18-month Workprogramme

September 2005 – February 2007

IST-004527 ARTIST2

Scientific Coordinator: Joseph Sifakis
Technical Coordinator: Bruno Bouyssounouse

Based on input from the Cluster Leaders:

Bengt Jonsson (Uppsala)

Giorgio Buttazzo (Scuola Superiore Sant'Anna - Pisa)

Reinhard Wilhelm (Saarland University)

Lothar Thiele (ETH Zurich)

Karl-Erik Arzen (Lund University)

Kim Larsen (Aalborg)

Table of Contents

1. Joint Programme of Activities

1.1 Overview

The NoE is structured and functions as a distributed laboratory – composed of its virtual teams (clusters) which perform research activities. These JPRA activities can be internal to the clusters ("Cluster Integration"), or between clusters ("NoE Integration") to achieve the overall integration of the NoE. These are detailed in the diagram on the following page.

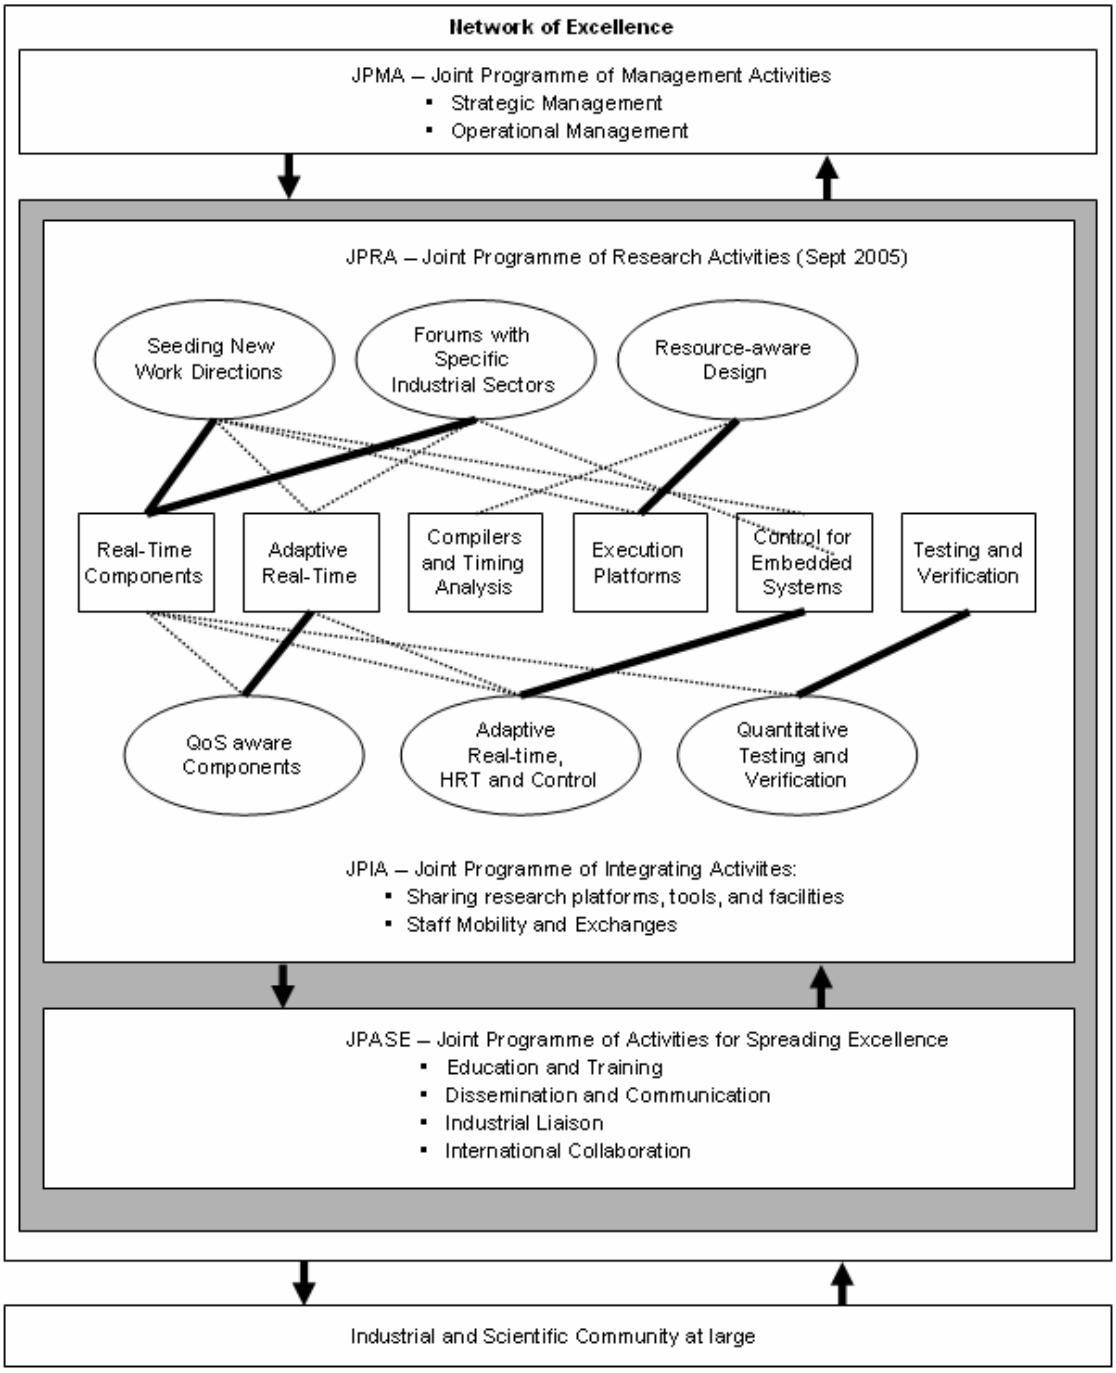

The “Cluster Integration” activities are managed within the cluster. For clarity they are not shown in the diagram, but are essential for integration at the cluster level. The “NoE Integration” activities - shown as ovals in the diagram – involve many clusters, and are driven by one cluster (identified by the thicker line).

The JPIA activities are transversal to the JPRA activities, and aim mainly to support their development through platforms, mobility, and infrastructure. They are strongly coordinated with the JPRA, and form the cement for the NoE’s internal integration.

The JPASE activities serve as the interface between the NoE and the community at large.

1.2 Proposed Changes at End of Year 2

In the coming year, we propose to implement the following changes to the Artist2 structure:

- Create a new activity in the Adaptive Real-Time cluster, called “Dynamic and Pervasive Networking”, led by Eduardo Tovar (Instituto Politécnico do Porto).
- Fusion of the ART Cluster-Integration activities: “Flexible Scheduling Technologies” and “Flexible Resource Management for Consumer Electronics”, to become “Flexible Resource Management for Consumer Electronics”. The merged activity will be led by Gerhard Fohler (TU Kaiserslautern).
- Incoming partner Sabine Glesner of TU Berlin plays an active part in the JPA implementation, and leads the Compiler Platform activity.
- ST Microelectronics becomes an affiliated partner.
- Continue to reinforce collaboration within the Compilers and Timing Analysis cluster. This can be implemented by launching joint events between the two communities. The activity “Architecture-Aware Compilation” is merged into the “Timing-Analysis Platform” activity, since these both related to the same underlying technical activity.

1.3 Governance

The governance structure is specified in the Consortium Agreement, and is reproduced below.

The methodology adopted for achieving the JPA objectives follows the same lines as for managing a laboratory. The activities, their objectives, their technical description, the partners involved, their roles, and the resources available have been clearly defined in the initial Description of Work, and updated in the deliverables. This is monitored and guided by a tight and rigorous management, as defined in the diagram below:

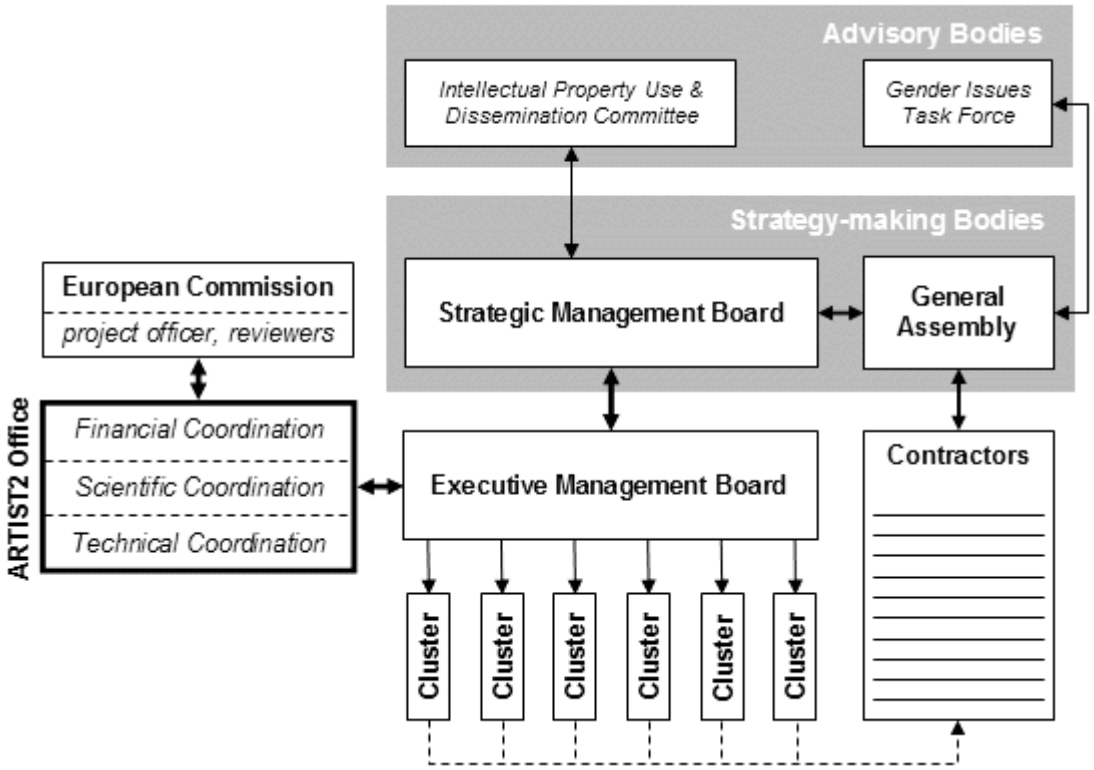

The main governance bodies are:

The **General Assembly** is composed of one representative per core partner. It is convened at the beginning of the project and meets once per year. It is chaired by the Scientific Manager.

The **Strategic Management Board** is initially composed of the NoE cluster leaders, and a representative of the Coordinator – who attends, with no voting rights. It is chaired by the Scientific Manager, assisted by the Technical Manager. It meets at least once per year – close to the General Assembly meeting. Its members are elected by the General Assembly every two years, according to modalities to be determined in the Consortium Agreement.

The **Cluster Leaders** (who compose the Executive Management Board) are responsible for the overall coordination of the activities led by their cluster. A cluster functions as a virtual team – with a degree of autonomy for defining its internal meetings and day to day management.

Over the course of the coming year, we do not intend to implement any changes to this management approach, but we do intend to reinforce, rationalize and insofar as is possible to automate the reporting. All the bodies have proven to be effective.

Within the consortium, we have acted to both improve and streamline the reporting procedures, and strengthen monitoring.

The NoE is moving forward with plans for setting up a sustainable structure for continuing interaction between research and industry, after the end of the NoE contract.

2. Project Timetable / Milestones

Given the size and complexity of the Artist2 NoE, it is preferable to use a detailed set of milestones to describe the JPA. This has the added advantage of presenting the past achievements and the expected evolution.

The JPA is organized into activities. The activities should not be considered as tasks of a workprogramme, with begin/end and synchronisation dependencies. Of course, the detailed description of an activity could be decomposed into sub-tasks and intermediate milestones, but this would imply a granularity that is too fine for research activities.

The inter-dependencies between activities are complex and rich, and will evolve dynamically. The work plan for the activities is provided in the 18 month workpackage descriptions.

The major milestones per activity are described below, updated at the start of Year 3:

2.1 Cluster: Real-Time Components

Cluster Leader: Bengt Jonsson (Uppsala)

The initial objectives of the different activities in the cluster were as follows.

- **Platform: Components Platform for Component Modelling and Verification**: to obtain initial versions of tool integrations for the component modelling and verification platform.
- **Cluster integration: Development of UML for Real-Time Embedded Systems**: to prepare an initial submission to the OMG standard for the UML profile for Modeling and Analysis of Real-Time and Embedded Systems (MARTE). This submission is approaching maturity, and is scheduled for completion at the end of 2006.
- **Cluster integration: Component Based Design of Heterogeneous Systems**: This is a new activity, replacing “Forums with specific industrial sectors” and “Seeding new research directions”. A description and the milestones are included in this document.

Milestones for the next 18 month period are as follows.

Component Modelling and Verification (Platform)
Susanne Graf (Verimag)

Milestones as foreseen in the last 18 month workplan and their realisation

- Year 1: Initial definitions of modules to assemble in the platform

  This milestone had been achieved at the end of year 1

- Year 2: Initial connections within a common framework of existing UML-based analysis and validation tools.

  This milestone has been achieved at the end of year 2: there exist new tool connections in the platform picture that can be demonstrated, including complete chains from modelling to validation, in particular

- *the Persiform tool chain from an Activity Diagram oriented UML profile for functional service specifications or annotated MSC to the SES workbench performance analysis tool.*
- *the Kermeta – IF tool chain manages software development support starting from the specification of components and their composition, their verification down to the generation of Java or C based executable units for the execution of a specific platform,*
- *the BIP/THINK tool chain represents the backend of a tool chain of a tool chain with the same motivations as the previous one.*
- *The start of the OpenEmbeDD, System@tic/Usine Logicielle, and SPEEDS project represent an important milestone, as their aims are fully in line with those of the platform and they provide the funding for deep technical work and the modelling languages they build upon, focus on different, complementary aspects.*

- Year 3: Strengthen and extend the existing tool chains so as being able to connect some of the analysis and validation tools developed by the partners or outside Artist to UML tools by means of mappings to a few semantic frameworks, in turn mapped to the input languages of the tools. This will allow realising tool chains from high level languages down to code.

  This work will include in particular, mappings from the HRC model defined in SPEEDS into semantics level formalisms for the connection to validation and and analysis tools as well as tools for model-based code generation.

- Year 4: Final integration of the results of the related Joint Research Activities.

Development of UML for Real-Time Embedded Systems (Cluster integration)

- One plans mainly two actions: -i- to organize a review of the Marte standard by Artist2 partners and conduct a specific workshop to debrief and finalize a common Artist2 review of the Marte standard. This action should be started at the beginning of December 2006. -ii- to conduct an industrial workshop (beginning of 2007) I order to disseminate the Marte standard and promote its usage.

Component Based Design of Heterogeneous Systems (Cluster integration)

- Year 3:
  - Unification of models of computation and comparison beween frameworks using denotational and operational semantics.
  - Rich heterogeneous interfaces and associated verification techniques based on Assume/Guarantee.
  - *A meeting on* Integrated Modular Avionics *and its impact of embedded systems design in avionics; will be scheduled during spring 2007; expected for fall 2007. The approach for this meeting will be similar to the one followed for the* Beyond Autosar *meeting.*
- Year 4:
  - Definition and classification of unified frameworks encompassing heterogeneity.

- Verification framework for rich heterogeneous interfaces.
- Organize a forum on the topic of design methods and tools for heterogeneous large embedded systems; deliver a summary of findings and recommendations for research on the topic of design methods and tools for heterogeneous large embedded systems.

2.2 Cluster: Adaptive Real-Time

Cluster Leader: Giorgio Buttazzo (Pisa)

A Common Infrastructure for Adaptive Real-Time Systems

Giorgio Buttazzo (Pisa)

- Year1 **(achieved)**: 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.*
- Year2 **(achieved)**: Deploy 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.*
- **Year3: Participate in the evolution of RTOS and networking standards, by introducing advanced scheduling methods for enhancing the predictability of real-time systems and handle their increased complexity.**
- **Year4: Identify the problems to be solved for developing a component-based real-time operating system.**

Flexible Resource Management

Gerhard Fohler (TU Kaiserslautern)

- Year 3:
  - Deduce QoS requirements from case studies (e.g., media processing) to provide operational parameters of the computing and communication infrastructures to allow the creation of global mechanisms for resource management.
  - Define architectural model for the framework for flexible scheduling that integrates multiple resources, including CPUs and networks: dynamically reconfigurable modules, multiple processors, interrupts with time protection, shared resources with time protection, memory protection, and energy/power-aware scheduling
- Year 4:
  - Provide a framework that allows the seamless integration of flexible scheduling schemes for integrated resources, allowing the choice of appropriate scheduling methods for individual activities in the different resources, and integrating application adaptation processes

QoS aware Components
Alejandro Alonso (UP Madrid)

- Year 1 **(achieved)**: Identification of the concrete integration topics: modeling of QoS properties in design models and components frameworks.

- Year 2 **(achieved)**: Study and dissemination of the approaches from different partners. Definition of case studies for comparing the approaches and begin its modeling. *The work has concentrated on UML profiles for the description of extra-functional properties and on evolutions of CCM and Robocop as the components frameworks.*

- **Year 3: Completion of the use cases using the different modeling approaches. Comparison and identification of guidelines on their use. Refinement of the modeling of some specific QoS properties and automatic model generation.**

- **Year 4: Propose a modeling techniques that combines the best features of both for some selected extra-functional properties. Propose requirements for future QoS support on components framework. Develop prototypes for proving the validity of some of the new identified new features.**

Real-Time Languages
Alan Burns (York)

Note that the activity on Real-Time Languages started in March 2006, so Year 3 milestones are after 18 months of activity:

Year2 **(achieved)**:

- Preliminary work in defining the future milestones and undertaking the necessary planning for future workshops, meetings, and joint work.

Year 3:

- **Organise and participate in the 13th IRTAW.**
- **Publish via a web site an initial set of patterns (repository) for use by Ada 2005 application programmers.**

Year 4.

- **Produce a white paper linking all real-time language work within ARTIST partners (including reference to external research effort where appropriate).**
- **Extend the repository**

Dynamic and Pervasive Networking Eduardo Tovar (Instituto Politécnico do Porto)

Note that the activity on Dynamic and Pervasive Networking is going to start on October 2006:

Year 3

- **Organise a kick-off cluster meeting on this activity within the 4th quarter of 2006.**
- **Produce a white paper on taxonomy of Wireless Sensor Networks (WSNs) and Mobile Ad-Hoc Networks (MANETs), elaborating on exemplificative applications, on their requirements and on how these map into technology design issues (1st quarter of 2007).**
- **Identify and characterize network protocols to support integrated and dynamic resource management in distributed environments as necessary for on-line adaptation and reconfiguration.**
- **Organise and participate in the 6th International Workshop on RTN (3rd quarter 2007).**
- **Concrete contributions on MAC and Routing protocols for WSN, MANETs, systems of embedded systems and adaptive distributed embedded systems.(4th quarter 2007).**

Year 4

- **Contributions on distributed computing paradigms (e.g., computation of aggregate quantities, collaborative computing, reconfigurable systems) as well as on dynamic QoS management, flexible scheduling and generally resource management in distributed systems exploiting previously proposed mechanisms (e.g., MAC and routing protocols).**
- **Organize a summer school on Real-Time Networks, involving key players from industry and academia, possibly focusing on specific topics such as WSN and MANETs.**
- **A SOTA report on WSN and MANETs, with web publishing.**
- **Contributions to the standardization bodies (e.g., IEEE 802.15.4, IEEE 802.11.x, IEEE 802 AVBridges).**

2.3 Cluster: Compilers and Timing Analysis

Cluster Leader: Reinhard Wilhelm (Saarland)

Timing Analysis Platform Reinhard Wilhelm (Saarland University)

- *Year2: Standard tool architecture and interfaces*
  The chosen interface language, AIR, is being extended by Saarland University and by AbsInt to suit the needs of other partners.

Four partners of the team (Vienna, Mälardalen, Tidorum, AbsInt) will continue to work on path description attributes for AIR to arrive at a uniform notation.

- *Year3: Initial integration of existing components*
  Mälardalen will wrap up its flow analysis into a component with well-defined interfaces, which will be integrated with the aiT tool of AbsInt and the Bound-T tool of Tidorum.
- *Year4: Version 2 integration of existing components*

Compilers Platform
Sabine Glesner (TU Berlin)

- **Year 1**: *Initial definition of common compiler platform*
  This milestone has been achieved by selection of ACE´s CoSy platform as the primary platform for most cluster partners.
- **Year 2**: *Initial implementation of the platform*
  This milestone has been achieved by installing and adopting the platform at the partners´ sites (partially after some setup meetings and training) for teaching and research purposes (e.g. for projects related to the architecture aware compilation activity). Examples: Aachen is using CoSy presently for development of SIMD and conditional instruction based code optimization in close cooperation with ACE. Likewise, TU Berlin (new core partner) is using CoSy for research on new compiler verification technologies. TU Vienna has integrated the Program Analyzer Generator (PAG) in the C++ infrastructure ROSE which is the source-to-source component of the compilers platform. 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.
- For **Year 3** and **Year 4**, the compiler cluster envisions a status where more and more new technologies (e.g. code optimization, verification) have been integrated into the common platform, which would lay a solid basis for self-sustained cooperations after the ARTIST2 funding period.

2.4 Cluster: Execution Platforms

Cluster Leader: Lothar Thiele (ETHZ)

WP1 Platform: System Modelling Infrastructure
Jan Madsen (Technical University of Denmark)

- Year2: Initial definition of the modelling platform (achieved). Several simulation- and formal-based models have been investigated and extended towards integration. Early integration of the simulation-based models, ARTS and MPARM, and of the formal-based models SymTA/S and Real-Time Calculus has been achieved. Initial linking between simulation- and formal-based models, MPARM and Real-Time Calculus has been investigated.
- Year3: Version 1 of the system modelling platform implementation. Due to the experience gained with the different modeling formalisms, it was found that rather

than aiming for a single unified model, the focus should be on further exploration of the existing models and in particular their interaction. Therefore, the focus will be on linking and integrating different modeling formalisms and to extend the models to support analysis and exploration as needed by the other cluster activities.

- Year4: Integration of modelling formalisms covering different levels of abstraction. Effective strategies for selecting model formalisms. Refinement and dessimination of models.

WP3 NoE Integration: Resource-aware Design

Luca Benini (University of Bologna) and Peter Marwedel (University Dortmund)

- Year2: A set of tools that can interact and work together and demonstrate the achievable optimizations on a particular hardware platform (achieved). Integration between AACHEN Lisa tools and Unibo's MPARM has been achieved. An early version of the memory aware compiuler by Dortmund has been targeted to the MPARM platforms.
- Year3: Strengthening the integration between Dortmund and Bologna: development of a memory-aware compiler for parallel multi-task applications. Linkoeping will also work to an integrated execution analysis environment for muilti-core systems.
- Year4: A methodology for the design of predictable embedded systems

Communication-centric systems (Cluster Integration)

Rolf Ernst (TU Braunschweig)

- Year1 (achieved): Assess the state-of-the-art in models that take into consideration the particularities of various quasi-standard communication protocols during system analysis and scheduling
- Year2: New best-case/worst-case models for hard real-time systems and at combined statistical and interval models for QoS applications in multi-media. These models may combine communication and computation, different models of computation, event models and scheduling policies (achieved)

  Due to the feedback of practical applications it was found that interval models cover a wide range of QoS application. Therefore, it was considered as more useful to extend work in this direction rather than starting a new work in the domain of combined statistical and interval models. This is subject to future work.

  The following main results were achieved:

  The industry increasingly applies hierarchical communication protocols, such as the automotive FlexRay Standard. New timing analysis techniques were developed to follow that trend. Since system dependability and flexibility are growing demands in embedded systems, techniques for fault tolerance and robustness measurement were developed and included in communication modelling and optimization. Additionally, case studies were performed to demonstrate the feasibility and practicability of the research results. Power optimization techniques were developed, since low power design is a new and urgent requirement in mobile and streaming applications.
- Year3: Analytic methods to estimate system properties
- Year4: Refinement and dissemination of these methods

Design for Low Power (Cluster Integration)
Luca Benini (University of Bologna)

- Year2: Component models will be investigated that model power dissipation of system components (achieved)) this objective has been achieved: several extensions to the MPARM platform power modelling capabilities have been developed (miulti-cluser systems, multi-frequency domains, multiple-voltage domains)
- Year3: Effective strategies for power management and power aware allocation and scheduling for both single-chip and distributed systems
- Year4: Integration of the different levels of abstraction - from scheduling via operating systems to system design - participating in low power design

2.5 Cluster: Control for Embedded Systems

Cluster Leader: Karl-Erik Arzen (Lund)

Control of Real-Time Computing Systems (Cluster Integration)
Karl-Erik Arzen (Lund)

Year1 Milestone: Roadmap describing the current state-of-the-art and the important research issues *(Achieved)*

The roadmap has been completed and partially disseminated. What remains is to make the entire roadmap more easily available to the general public. We plan to print the two roadmaps together as a single document or report during the fall of 2006.

Year2-4 Milestones:

- Progress made on the fundamental underlying issues: decreased requirements on prior knowledge about resource utilization, increased possibilities to use COTS implementation platforms, and enhanced robustness towards load variations (*Achieved to 30 % currently)*

The research performed during this year contribute to the solution of several of the above items. For example, the work on feedback control of Linux scheduling is a step towards being able to utilize COTS implementation platforms, and the work on queueing system models is motivated by the aim to be robust against load variations.

New Year 3-4 Milestone:

→ **Increase our international and industrial visibility. A good means for this is through the organization of and the participation in the FeBID workshops.**

Real-Time Techniques in Control System Implementation
(Cluster Integration)
Alfons Crespo (Universidad Politécncia de Valencia)

Year1: Roadmap describing the current state-of-the-art and the important research issues (*Achieved*)

Year2: A common framework of the control parameters that can be influenced by an embedded control system implementation and the real time operating systems criteria that can be adjusted to increase the robustness of the control system (*Achieved to 50%*)

This milestone has not been fully completed yet. Our aim is to complete this during the reminder of 2006

Updated Year3-4 milestones:

- → **A common framework model in order to facilitate the control and computing co-design**
- → **Organization of an annual Graduate School on Embedded Control Systems**
- → **Organization of a follow-up of the Lund Workshop on Control for Embedded Systems**

Adaptive RT, HRT and Control (NoE Integration)
Karl-Erik Arzen (Lund)

Year1: Setting the technical background and assess the needs *(Achieved 100%)*

Year2: Demonstrate that applications of diverse type can be specified in terms of resource-aware tasks (*Achieved 80 %)*

The work within the activity has focused on two application types only: multimedia applications and real-time control. Within these two broad application areas, several types of application have, however, been studied. These two application types are also the ones that are most natural for these techniques.

Update Milestone for Year3:

- — **Demonstrate that scheduling algorithms can be made adaptive by means of control schemes**
- — **The organization of a new industrial workshop along the lines of the workshop organized jointly with the Beyond AUTOSAR activity**
- — **The organization of a follow-up research workshop to the Lund Workshop on Control for Embedded Systems held in June 2005. The workshop is currently planned for Jan-Feb 2007**

Design Tools for Embedded Control (Platform)

Martin Törngren (KTH)

Existing milestones - Year1-2: Identification of which of the existing tools that will be included in the platform, and specification of their interfaces

Comment: The tools developed by the cluster have been investigated and compared. Functionalities represented by other discipline's tools have also been investigated. Interfaces have been described at a high level of functionality. Different approaches to model and tool integration have been investigated. The individual tools have been further developed and disseminated. One prototype tool integration platform has been developed.

Exisiting milestone - Year3: Develop the necessary interfaces that allow the individual tools to be used together

- **Development of integration scenarios**
- **Performed several case studies on model and tool integration, involving tools specific to the cluster as well tools typically dealt with by other research communities (clusters)**

Comment: As a basis for tool integration, it is important to clarity the relevant usage scenarios, i.e. how the tool integration supports the various design activities. Moreover, systems design is not limited to just the aspects traditionally dealt with by this cluster. Therefore it is important to carry out case studies that illustrate tool integration also considering other relevant aspects. This update of the Year 3 milestone is also supported by the previous 18 month plan..

Existing milestone - Year4: Usage of the tools in new co-design based research activities, adoption in industrial case studies.

2.6 Cluster: Testing and Verification

Cluster Leader: Kim Larsen (Aalborg)

Testing and Verification Platform for Embedded Systems (Platform)

Kim G. Larsen (BRICS/Aalborg)

- Year2: A server on which the main testing and verification tools developed and used by the participants will be installed and configured.
  Design of a coordination layer for parallel and distributed model checking, Design of a GRID infrastructure, links to mature model checking tools via the Yahoda homepage.
- Year 3:
  - Links to the tools developed and applied by the partners will be collected at a common web entry. Also, it will be analysed whether a common web interface can be provided for tool invocation in a trusted and controlled manner. This is a revision of the above milestone on providing a single powerful server for all tools.
  - The ongoing work on tool evaluation through case studies will be continued and made accessible at the open repository. Also, links to mature version swill be provided via the Yahoda tool homepage. This is a revision of the above milestone on links to mature versions.
  - Further experiments on exploiting contemporary technologies (GRID and PC clusters) will be made. This includes experiments on establishing tool access on available sites (e.g. NorduGrid) as well as further development of distributed model checkers.
- Year4: Integration of results from the related Joint Research Activities

Quantitative Testing and Verification (Cluster Integration)

Ed Brinksma (University of Twente)

- (achieved) Year1: Initial results for testing and verification with emphasis on quantitative aspects
- Year2: Develop theory, methods and tools for testing and verification of embedded systems with emphasis on quantitative aspects (e.g. real-time and stochastic phenomena) that are of particular importance for the correctness of embedded systems.
  Further work on robustness, metrics and abstraction. Also collect and classify major case studies.
- Year3:
  - Development of algorithms and implementation of tools for optimal controller synthesis, robust model checking, coverage-based test selection and code generation.
  - Existing verification tools and test generation tools are more strongly connected, including stronger links between academic and industrial tools.
- Year4:
  - Development of generic framework using abstraction and compositionality for efficient analysis of quantitative models.

- Emergence of a range of new powerful debugging and analysis based on various combinations of testing and verification techniques.

Verification of Security Properties (Cluster Integration)
Sandro Etalle (Twente)

- (achieved) Year1: Define a reference model for security protocols
- Year2: prototypes capable of performing automatic analysis of security protocols
- **Year 3: Development compositional proof techniques for verifying services security properties, and for verifying group protocols.**
- Year4: Design monitoring procedures for ensuring trust in services execution**.**

2.7 Global NoE Activities

JPIA-Staff Mobility and Exchanges (WP1)

- **WP1: Staff Mobility and Exchanges**
  All partners
  - In years 1&2, staff Mobility and exchanges were a successful component of integration. (achieved)
  - **Over the course of the next 18 months, we expect to have the same level of staff mobility as was the case for Years 1&2 (described in sections 4-9).**
    In particular, UJF/Verimag and other partners will fund travel for up to 10 visiting researchers from various parts of the world, to participate in events organised by Artist2, which may be held anywhere in the world.

JPASE-Education (WP2)

- **WP2: Courseware**
  - Clearly, the way forward for courseware is a bottom-up approach, in which we provide links to existing high-quality course materials and contacts with teachers. These are made a vailable to the general public via the Artist web portal. (partially achieved)
  - **Collection and dissemination of courseware will continue and be extended within Years 3 and 4.**

- **WP2: Support for Summer Schools, workshops, and conferences**
  Bruno Bouyssounouse (Verimag)
  The complete list of schools organized by or in collaboration with Artist is available here: http://www.artist-embedded.org/artist/-Schools-.html
  UJF/Verimag may provide up to 25k€ to various events in the area, through direct sponsoring, by handling travel/accommodation fees for speakers, or other means as appropriate.
    - Past Schools organized by or in collaboration with Artist2 is very impressive (achieved).
      These include:

      ADSD 2006: Advanced Digital Systems Design
      September 25-29, 2006 Lausanne, Switzerland

      FOSAD 2006: 6th International School on Foundations of Security Analysis and Design *September 10-16, 2006 Bertinoro, Italy*

      MDD4DRES : MDE Approaches
      September 4-8, 2006

      First European ARTIST Laboratory on Real-Time and Control for Embedded Systems *July 10-14, 2006 Pisa, Italy*

      ARTIST2 / UNU-IIST Spring School in China 2006
      April 3-15, 2006 Xi'an, China

      ARTIST2 Graduate Course on Embedded Control Systems
      April 3-7, 2006 Prague, Czech Repbulic

      ARTIST2 Summer School 2005
      Component & Modelling, Testing & Verification, and Statical Analysis of ES
      September 29th - October 2nd 2005 Nässlingen, Sweden
      ARTIST2 Summer School on

    - **Two Artist2 schools are planned within Year 3:**

      ARTIST2 - MOTIVES 2007
      (http://www.artist-embedded.org/artist/-MOTIVES-2007-.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.

      Artist2 / UNU-IIST School in China - 2007
      (http://www.artist-embedded.org/artist/-Artist2-UNU-IIST-School-in-China-.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).

    - **Further schools will be organised in Year 4 as well.**

JPASE – Dissemination and Communication (WP2)

- **WP2: Organisation and Support for Conferences, Workshops, Seminars**
  Bruno Bouyssounouse (Verimag)
    - In years 1&2, Artist2 has organized or participated in the organisation of a very large number of workshops in the area. (achieved)
      The full set of workshops, including resulting materials are made available to

the general public, via the Artist web portal, here:
http://www.artist-embedded.org/artist/-Workshops-and-Seminars,29-.html

- **In Year3, Artist2 will directly organize the following workshops and seminars (see same link as above for details):**

  MARTES 2006 ***October 2nd, 2006 Genova, Italy***
  This workshop gathers researchers and industrial practitioners to survey modeling and model-based analysis of distributed, real-time and embedded systems.

  JTRES 2006 ***October 11-13, 2006 Paris, France***
  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.

  Foundations and Applications of Component-based Design
  October 26th, 2006 Seoul, South Korea
  The workshop aims to gather together researchers from computer science and electrical engineering and will seek a synthesis between the underlying paradigms and techniques.

  WESE'06 - Embedded Systems Education
  October 26th, 2006 Seoul, Korea
  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.

  MoCC - Models of Computation and Communication
  November 16-17, 2006 Zurich, Switzerland
  Communication and cooperation between several disciplines: software and hardware but also computer science and engineering, real-time and distributed systems, telecommunication, control and signal processing.

  Timing Analysis
  November 17th, 2006 Paphos, Cyprus
  1-day workshop. This Special Track will be concerned with questions around the integration of timing analysis in the industrial development process.

  ARTIST2 Workshop on Basic Concepts in Mobile Embedded Systems
  December 3-4, 2006 Vienna – Austria
  It is the objective of this workshop to elaborate the basic concepts on mobile embedded systems based on existing approaches in distributed, real-time, and dependable systems.

- **In Year3, Artist2 will also participate in the organization of the following workshops and seminars (see same link as above for details):**

  ARCS 2007 ***March 12-15, 2007 Zurich, Switzerland***
  ARCS 2007 will cover a broad range of research topics related to basic technology, architecture, and application of computing systems. Organic Computing (Autonomic or Proactive Computing) may help to manage the increasing complexity of computing systems. ARCS 2007 is intended for researchers, R&D engineers, and practitioners.

SCOPES 2007
April 20th, 2007 Acropolis, Nice, France
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.

- **WP2: Web Portal for Dissemination**
  Bruno Bouyssounouse (Verimag)
  - Over the course of the second year, a web portal as described in the Description of Work has been set up. Although this was later in the year than initially planned, the portal is arguably far more complete than the initial version initially described. We believe that it already serves as a reference for the entire embedded systems community (research & industry).
    (achieved)
  - Statistics on web accesses and mail addresses registered with the site are be collected and summarised on a monthly basis within year2 for inclusion in the next management deliverables (see Spreading Excellence).
    (achieved)
  - The requirements and process for joining the network as an affiliate will be simple and easily accessible on the web portal. All ARTIST2-affiliated authors are encouraged to indicate this affiliation in published papers.
    http://www.artist-embedded.org/artist/Becoming-an-Affiliated-Partner.html
    (achieved)
  - The web portal will include the following elements relevant for Embedded Systems Design:
    - Documentation on the state-of-the-art upon which the NoE is basing its work. This state-of-the-art collection could include the various roadmap documents that have been produced in ARTIST and ARTIST2, or could be extracted from the existing deliverables.
      http://www.artist-embedded.org/artist/State-of-the-Art.html
      (achieved)
    - Access to main documents, including books, journals, position papers, documentation for standards, newsletters, and in particular, documents produced by Artist2. (achieved)
      http://www.artist-embedded.org/artist/-Artist2-Dissemination-.html
    - Information about main events, including conferences, schools, seminars, high-level events, and in particular the events organised by Artist2. http://www.artist-embedded.org/artist/-Home-Page-.html
      (achieved)
    - Access to courseware collected by the Artist2 partners from education events, as well as links to other sites providing materials.
      http://www.artist-embedded.org/artist/-Course-Materials-.html
      (achieved – continuing into Year 3)

- Management of a set of mailing lists, for selected themes (upcoming thesis presentations, open positions, calls for papers, etc).
  (not yet achieved – planned for Year 3)

- **WP2: Newsletter**
  Bruno Bouyssounouse (Verimag)
  - The NoE will publish a Newsletter providing information about important events for research on Embedded Systems Design. The format and distribution list for the newsletter will evolve over time.
    http://www.artist-embedded.org/artist/Artist2-Newsletter,438.html
    (achieved)

    Topics may include the following, when these are both available (eg: authorised by the original authors), and of sufficient interest to the general community:
    - *Events organised by Artist*
    - *External events*
    - *Technical overviews of interest to the general community*
    - *Interviews from leading research or industrial figures*
    - *A word from the project officer*

JPASE – Industrial Liaison (WP2)

- **WP2: Links to Industry**
  Bruno Bouyssounouse (Verimag)
  - High-level participation with industry:
    Actions for structuring industrial relations between ARTIST2 and European R&D, through involvement in Integrated Projects in the governance of the ARTEMIS European Technology Platform, and in the accompanying ARTEMISIA association.
    Artist2 has actively participated in writing the ARTEMIS Strategic Research Agenda, and will continue to contribute to its evolution in Year 3.
  - Deep, individual ties to industry
    Artist2 has over 26 affiliated industrial and SME partners – and this is reflected in the strong ties to industry by clusters, activities and individual partners, described in the Activity deliverables.
    The full set of Affiliated partners is shown on the Artist2 web portal:
    http://www.artist-embedded.org/artist/-Affiliated-Partners-.html
    (achieved)

- **WP2: Contributions to Standards**
  Bruno Bouyssounouse (Verimag)
  - Year1-4: In Year3, we will continue ongoing work on standards.

    The list of ongoing actions is described in the deliverable on Spreading Excellence, as well as the Artist web portal:
    http://www.artist-embedded.org/artist/-Standards,70-.html
    (achieved, ongoing)

JPASE – International Collaboration (WP2)

- **WP2: International Collaboration**
  Bruno Bouyssounouse (Verimag)
  http://www.artist-embedded.org/artist/-International-Collaboration-.html
  International Collaboration events are intended to gather together the very best world-leading experts from academia and industry, to discuss progress on the state of the art, relevant work directions. The annual events are each focused on a specific topic, and include presentations from leading figures in Europe, the USA, and Asia.
  - In Years 1&2, the following International Collaboration events were organized :

    ARTIST2 Spring School in China 2006
    April 3rd – 15th 2006 Xi'an, China
    The first ARTIST / UNU-IIST Spring School has been held in Xi'an, China, April 3rd – 15th 2006, and gathered more than 50 participants, of which approximately 40 were students from the top universities in mainland China. Given the success of this first edition, it has been decided to organise a second ARTIST2 school in China, near Shanghai in 2007.

    Joint US-EU-TEKES workshop
    June 21-22 2006 Helsinki, Finland
    Workshop held under the auspices of NSF, the EU's IST Program and Tekes, the Science and Technology Agency of Finland.

  - In Year 3, Artist2 will organize the following International Collaboration event :

    Artist2 / UNU-IIST School in China – 2007
    Suzhou (near Shanghai), August 1-10 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. Artist2 lecturers include Karl-Erik Arzen (Lund), Luca Benini (Bologna), Paul Caspi (Verimag), and Kim Larsen (Aalborg).

2.8 Provisional Budget: Sept 2005 – February 2007

4 Year Total: 6 500 000 € 18 Month Total: 2 437 500 €

			18 month amounts:
JPMA	*WP0*	**Joint Programme of Management Activities**	**170 625 €**
		Artist2 Office	170 625 €
JPIA	*WP1*	**Joint Programme of Integration Activities**	
		Sharing Research Platforms, Tools Facilities	**527 388 €**
		Platform for Component Modelling and Verification	73 128 €
		A common infrastructure for adaptive Real-time Systems (Platform)+E126	57 689 €
		Timing - Analysis Platform (Platform)	88 841 €
		Compilers Platform (Platform)	87 750 €
		System modelling infrastructure (Platform)	70 608 €
		Design Tools for Embedded Control (Platform)	65 004 €
		Testing and Verification Platform for Embedded Systems (Platform)	84 369 €
		Mobility	**268 704 €**
		Mobility between partners' sites (core or not)	268 704 €
JPASE	*WP2*	**Joint Programme of Activities for Spreading Excellence**	**292 500 €**
		Education (Summer & Graduate Schools, courseware)	73 125 €
		Dissemination and Communication (workshops, seminars, web portal)	146 250 €
		Industrial Liaison (direct interaction with leading companies)	48 750 €
		International Collaboration (High Level Events)	24 375 €
JPRA		**Joint Programme of Research Activities**	
	WP3	NoE Integration	**574 155 €**
		Forums with specific industrial sectors (NoE Integration)	167 211 €
		Seeding New Work Directions (NoE Integration)	83 606 €
		QoS aware Components (NoE Integration)	42 377 €
		Resource-aware Design (NoE Integration)	119 435 €
		Adaptive Real-time, HRT and Control (NoE Integration)	82 122 €
		Quantitative Testing and Verification (NoE Integration)	79 406 €
	WP4	*Merged into WP5 (when Components and Modelling merged with HRT to form Real Time Components)*	
	WP5	Cluster Integration for cluster: "Real-Time Components"	**37 782 €**
		Development of UML for Real-time Embedded Systems (Cluster Integration)	37 782 €
	WP6	Cluster Integration for cluster: "Adaptive Real-Time"	**176 240 €**
		Flexible Scheduling Technologies (Cluster Integration)	63 590 €
		Flexible Resource Management for Consumer Electronics (Cluster Integration)	59 049 €
		Real-Time Languages (Cluster Integration)	53 601 €
	WP7	Cluster Integration for cluster: "Compilers and Timing Analysis"	**55 575 €**
		Architecture-aware compilation (Cluster Integration)	55 575 €
	WP8	Cluster Integration for cluster: "Execution Platforms"	**153 825 €**
		Communication-centric systems (Cluster Integration)	80 694 €

	Design for low power (Cluster Integration)		73 131 €
WP9	Cluster Integration for cluster: "Control for Embedded"	**122 504 €**	
	Control in real-time computing (Cluster Integration)		52 502 €
	Real-time techniques in control system implementations (Cluster Integration)		70 002 €
WP10	Cluster Integration for cluster: "Testing and Verification"	**58 203 €**	
	Verification of Security Properties (Cluster Integration)		58 203 €
	Totals:	*2 437 500 €*	*2 437 500 €*

2.9 Indicative Efforts (Sept 2005 – February 2008)

	JPMA	JPIA	JPASE	JPRA								
	WP0	WP1	WP2	WP3	WP4	WP5	WP6	WP7	WP8	WP9	WP10	
*Figures are in man*months*	**Management**	**Platforms & Mobility**	**Spreading Excellence**	**NoE Integration**	*Modelling and Components cluster halted*	**Real-Time Components**	**Adaptive Real-Time**	**Compilers and Timing Analysis**	**Execution Platforms**	**Control for Embedded**	**Testing and Verification**	**TOTALS**
CDC	26											26
UJF/Verimag	26	5	3	14		2					2	52
Aachen		7	12	4								23
Aalborg		5	1	2							2	10
Absint		7	1									8
Aveiro		2	1				4					7
Cantabria		1	2				4					7
CEA		4	1	3		4						12
CFV		2	1	2							2	7
Czech		5	1	2						4		12
Dortmund		8	1	4								13
DTU		4	2	3					5			14
ETHZ		1	1	5					6			13
FTRD			1	2							2	5
INRIA		7	2	9		1						19
KTH		6	3	2						7		18
Linkoping		5	1	3					4			13
CNRS		2	1	2							2	7
Lund		5	1	2						7		15
Malardalen		3	4									7
OFFIS			1	5								6
PARADES			1	7							1	9
Pavia												0
Madrid		2	1	4			4					11
Saarland		8	1	4								13
ST												0
Eindhoven			1						6			7
Vienna		6	1	9								16
TUBS		5	2						5			12
Twente		2	1	4							2	9
Bologna		4	3	3					5			15
Uppsala		12	2	6								20
UPVLC		5	2	2						7		16
York		4	1				5					10
Porto		1	4				4					9
EPFL		2	1	9								12
Pisa		8	5	2			11					26
ACE		5	3									8
Tidorum		3	1									4
KaiserSlautern		3	1	2			6					12
Berlin		6	1									7
TOTALS	52	155	73	116	0	7	38	0	31	25	13	510

In this table,

- INRIA refers to INRIA/IRISA/Université de Rennes 1
- CNRS refers to CNRS/ENS Cachan/LSV
- UJF/Verimag refers to UJF/Verimag, INPG/Verimag and CNRS/Verimag

Please note that WP7=0, because the Compilers&Timing Analysis activities are now all platforms (WP1).

2.10 List of Workpackages and Deliverables

The description of the Joint Programme of Activities for the coming 18 months is composed of the following workpackages. Each workpackage consists of a set of activities.

Joint Programme of Activities
(18 months period Sept 2006 – Feb 2008)

All deliverables listed here are: of nature "R"=Report, Dissemination Level = "PU"=Public, and to be delivered at project month 36 (Sept 2007).

WP	Work package title	Lead contractor	Start month	End month	Deliverable ID
WP0	JPMA : Joint Programme of Management Activities	1 CDC	0	48	**D1-Mgt-Y3 Year 3 Project Management Report**
		2 UJF/ VERIMAG	0	48	**D2-Mgt-Y3 Year3 Project Activity Report**
WP1	JPIA : Joint Programme of Integrating Activities	2 UJF/ VERIMAG	0	48	**D4-RTC-Y3 Component Modelling and Verification (Platform)**
		37 Scuola Sant'Ana	0	48	**D7-ART-Y3 A common infrastructure for adaptive Real-time Systems (Platform)**
		25 Saarland	0	48	**D12-CTA-Y3 Timing - Analysis (Platform)**
		3 Aachen	0	48	**D13-CTA-Y3 Compilers (Platform)**
		12 DTU	0	48	**D14-EP-Y3 System modelling infrastructure (Platform)**
		16 KTH	0	48	**D18-Control-Y3 Design Tools for Embedded Control (Platform)**
		4 Aalborg	0	48	**D22-TV-Y3 Testing and Verification Platform for Embedded Systems (Platform)**
WP2	JPASE : Spreading Excellence	2 UJF/ VERIMAG	0	48	**D3-Mgt-Y3 Report on Spreading Excellence**
WP3	JPRA : NoE Integration - Research Activities	24 UP Madrid	0	48	**D8-ART-Y3 QoS aware Components (NoE Integration)**
		31 Bologna	0	48	**D15-EP-Y3 Resource-aware Design (NoE Integration)**
		19 Lund	0	48	**D19-Control-Y3 Adaptive Real-time, HRT and Control (NoE Integration)**

		30	Twente	0	48	**D23-TV-Y3 Quantitative Testing and Verification (NoE Integration)**
Please note that workpackages WP5-WP10 concern only Cluster integration (not NoE Integration), and do not include the Platforms (which are in WP1). **Workpackage 4 (Modelling and Components) was merged with HRT at the end of Year 1, to form the Real Time Components (WP5) cluster.**						
WP5	JPRA : Real-Time Components	8	CEA	0	48	**D5-RTC-Y3 Development of UML for Real-time Embedded Systems (Cluster Integration)**
		32	Uppsala	25	48	**D6-RTC-Y3 Component Based Design of Heterogeneous Systems (Cluster Integration)**
WP6	JPRA : Adaptive Real-time	7	Kaiser-slautern	24	48	**D9-ART-Y3 Flexible Resource Management (Cluster Integration)** *fusion of "Flexible Scheduling" and "Adaptive Resource Management for Consumer Electronics"*
		34	York	18	48	**D10-ART-Y3 Real-Time Languages (Cluster Integration)**
		35	Porto	25	48	**D11-ART-Y3 Dynamic and Pervasive Networking (Cluster Integration)**
The activity "Architecture-aware compilation (Cluster Integration)" was merged into the Timing-Analysis Platform at the end of Year2. This was the only activity in WP7.						
WP8	JPRA : Execution Platforms	29	TUBS	0	48	**D16-EP-Y3 Communication-centric systems (Cluster Integration)**
		31	Bologna	0	48	**D17-EP-Y3 Design for low power (Cluster Integration)**
WP9	JPRA : Control for Embedded Systems	19	Lund	0	48	**D20-Control-Y3 Control in real-time computing (Cluster Integration)**
		33	UPVLC	0	48	**D21-Control-Y3 Real-time techniques in control system implementations (Cluster Integration)**
WP10	JPRA : Testing and Verification	30	Twente	0	48	**D24-TV-Y3 Verification of Security Properties (Cluster Integration)**

3. Cluster: Real-time Components

Cluster Leader: Bengt Jonsson (Uppsala)

The following is a description of the activities and overall objectives for the period: September 2006 – February 2008. The next reporting period will cover September 2006 – August 2007.

Staff Mobility

Area of Collaboration:	**Generation of component models**
Cluster:	Real Time Components
Sending Partner:	Dortmund (affiliated)
Receiving Partner:	Uppsala (core)
Person:	Harald, Raffelt, M.Sc.
Technical Work:	Further development of the LearnLib Tool, plus preparation for major case study (telecommunications protocol)
Dates:	July 5 - Aug. 24
Costs:	Housing 2 kE, Travel (several trips) 2 kE, Misc. 1 kE

Past Meetings in Year 2

Distributed Embedded Systems", organized at the Lorentz Center in Leiden,
Nov. 21-24, 2005, by Ed Deprettere and Lothar Thiele
http://www.lc.leidenuniv.nl/lc/web/2005/177/info.php3?wsid=177

Meeting Beyond AUTOSAR, co-organized by Werner Damm (OFFIS) and Albert Benveniste (INRIA), 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

ARTIST2 Summer School 2005
http://www.artist-embedded.org/FP6/ARTIST2Events/SummerSchools/Artist05.html
September 29 - October 2, 2005, on Component Modelling, Testing and Verification, and Static analysis of embedded systems held at Nässlingen, Sweden, September 29 - October 2, 2005
The ARTIST2 Summer School was a 4 day summer school for young researchers working or wanting to work in the fields of modelling, validation and performance analysis of embedded systems as well as engineers from industry with practical background in design and testing of embedded systems.
It attracted more than 60 students from Europe, and was highly appreciated by attendants.

FORMATS 05 http://www.it.uu.se/formats05/
The 3rd FORMATS was held in Uppsala, Sweden, September 26 - 28, 2005 in conjunction with ARTIST2 summer school,

Scandinavian ARTIST2 Seminar on Embedded Systems Design, Aug. 21, Stockholm, Sweden, (click from www.snart.org)
organized by SNART (the Swedish National Real-Time Association www.snart.org)
This day consisted of presentations of the ARTIST2 NoE to Swedish researchers, students, industrialists, with the aim to discuss how to improve collaboration between research groups and between academia and industry.

Meetings Planned in Year 3

Models of Computation (Zurich, Paul Caspi)
Mobile Embedded Systems (driven by Vienna)
Dagstuhl seminar (March 4 - 9 1007) organized by Lothar Thiele
Forum with aeronautics sector on Integrated Modular Avionics (to be organized by Albert Benveniste)
Winter School 07 (see the ARTIST2 Newsletter)

Meetings Planned in Year 4

Meeting on predictability of hardware in automotive/avionics and semiconductor industry (to be organized by Werner Damm, Rolf Ernst, and Reinhard Wilhelm)

3.1 Platform: Components Platform for Component Modelling and Verification

Activity leader: Susanne Graf (Verimag)

3.1.1 Year 1 Achievements: Sept 2004 – August 2005

The main objective of the first year of the project was to obtain an inventory of potentially interesting work on tools, to do some 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.

Due to the large span of applications covered by the tools to be integrated into to the platform, this integration is not intended to be a strong integration in the classical sense of an integrated toolset, but rather a set of components that can be used in combination with specific components to form different tool chains. The baseline of the tools is that they are UML compatible and share subsets of UML profiles. Some components will be specific to particular tool-chains and whereas others are useful in several ones.

Presently considered sub platforms are identified by the following working titles:

- A platform for the development of safety-critical embedded systems
- A platform for the analysis of performance critical service-based systems
- A platform for the certification of smart-card applications

The relevant subsets of UML used in the context of these three environments are specific for the concerned target application types, at the level of requirements and abstract design which is the main focus of the platform in a first phase. We intend to share some of the analysis tools amongst the platforms thanks to the mapping into a common semantic level model. Also the profile concerning architecture modelling may be shared, but will be considered later.

All the tools developed are or will be ported to Eclipse.

3.1.2 Year 2 Achievements: Sept 2005 – August 2006

A detailed description of these achievements is provided in the activity's deliverable.

1. Modelling languages and interaction with standards

The work on the platform interacts with and depends on several activities related to the development of UML-based modelling languages and the development of formalisms for a semantic level representation of models. An important goal 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. This should finally allow handle models more than one profile simultaneously. This section presents both user-level and semantic-level formalisms. In fact, the separation between these levels is sometimes narrow, as user-level profile will lift some of the concepts useful for validation at the user level and in some cases may be used for both purposes. We start with those clearly meant as user level language and pass then to semantic level formalisms.

The **MARTE UML profile** [EDG+05] for modelling real-time systems and their non functional properties, plays a central role as a user level format. The effort involves *CEA,* INRIA, *Cantabria,* and *Carleton University* Canada (Dorina Petriu and Murray Woodside), with significant feedbacks from *INRIA* and *VERIMAG.* The work has well progressed in 2006, both on the general analysis profile [EMDG06] and for schedulability related issues [LMD06, MLD06]. It is now close to an acceptance as an OMG standard (see report on the standardization issues) and has started to be implemented. It will be a base for an open source modelling and validation platform developed in the OpenEmBeDD project. The development of this profile, its implementation as an Eclipse component and its integration to the RSA IBM tool requires a huge amount of work and feedback from a large panel of end users. For that, it benefits from the support of two large French projects:

- The Usine Logicielle (Software Factory) project of the System@tic pole of competitivity (involving as end users in addition to the previous partners: EADS, Dassault Aviation, Hispano Suiza, EdF, MBDA, CS… see www.usine-logicielle.org). In this project, MARTE standard is also implemented as an Eclipse component
- The OpenEmbeDD platform (involving as end users in addition to the previous partners: France Telecom, CS, Airbus)

These two projects also support the development of an action language editor (Eclipse component) to instantiate the UML action semantics on domain usage (syntax and refined semantics).

The work on MARTE is completed for the automotive domain by the development of the "EAST-ADL 2" UML **profile for automotive architecture and component modelling.** Based on the *Autosar*™ meta-model, it aims to provide a higher level of software component modelling and to better support behavioural modelling aspects. It complements for the automotive domain the MARTE profile, in particular, through specialisation of the UML component concept by a tight mapping on the Autosar™ paradigms. This Autosar™ mapping limits its capabilities of managing dynamic interaction among components and will benefit from on going works done in the cluster on "rich component models". It will provide an Eclipse component within the IST ATESST project developed by CEA, KTH and TUB together with their industrial partners (Volvo Tech., Daimler Chrysler, Siemens VDO, etc., see www.atesst.org).

Within the OPRAIL project, a UML profile called **Safe-UML** to be used in the context of safety critical system is being developed. The experiences with Safe-UML will be used within SPEEDS to derive efficient analysis techniques for UML/SysML. Safe-UML is a restriction of general UML to be used for enabling a CENELEC-conformant development of safety-critical rail systems. As UML is intended to cover the entire design process and when deployed in a particular domain of application, has to be instantiated for a concrete, tool-supported environment; in addition, different restrictions and specializations apply to different development stages. The profile focuses on structural diagrams (class diagrams) and behavioural diagrams (state charts). The main general restrictions concern event handling (no unbounded message queues) and control of non determinism (in concrete implementation specs). Part of the specifications concern UML on a general level and define Safe-UML (S), where (S) stands for seamless development. On a more concrete level, the UML tool Rhapsody in C++ is considered, yielding Safe-UML (R), (R) for Rhapsody. Finally, Safe-UML (V) treats formal verification with the model checker RUVE. The aim of the Safe-UML definition is twofold. On the one hand models following this profile shall be compliance to standards (for example code compliance with the German railway guidelines MÜ8004 for the generated code) and on the other hand it is expected that verification tools based on Safe-UML can significantly be improved from a performance point of view in relation to a general UML verifier.

The work to develop the concept of **rich component models** into a mature framework for system design is pursued within IP-SPEEDS by RTC partners INRIA, OFFIS, PARADES, and VERIMAG. They are currently developing a meta-model for rich components, called **HRC**. 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) [DVMJ05, Da06]. The work in SPEEDS also involves a new theory of *interfaces* is being developed, allowing for cross-viewpoint assume-guarantee reasoning. This piece of work undertaken within the SPEEDS project is a clear by-product of previous and current work developed in the ARTIST community (for more background, see RTC cluster report).

The **BIP framework** (Behaviour, Interaction, Priority) developed at VERIMAG over the last 5 years [GS05, BBS06] will play an important role in OpenEmbeDD, SPEEDS and other projects being set up for providing a mapping from user level languages to the semantic level, preserving the structure. It addresses two fundamental sources of heterogeneity: one is the composition of subsystems with different execution and interaction semantics. The second is the use of models that represent a system at different degrees of detail and are related to each other in an abstraction (or equivalently, refinement) hierarchy. A key abstraction in system design is the one relating application software to its implementation on a given platform. Application software is often largely untimed, whereas the application code running on a given platform, however, is a dynamic system that can be modelled by a set of timed or hybrid automata. The run-time state includes not only the variables of the application software, but also all variables that are needed to characterize its dynamic behaviour, such as time variables and other quantities used to model resources.

The aim of the BIP framework is to provide a semantic framework for such systems of heterogeneous components.

- It supports a component construction methodology based on the thesis that components are obtained as the superposition of three layers. The lower layer describes *behaviour*. The intermediate layer includes a set of *connectors* describing the *interactions* between transitions of the behaviour. The upper layer is a set of *priority rules* describing scheduling policies for interactions. Layering implies a clear separation between *behaviour* and *structure* (connectors and priority rules).
- It uses a parameterized *composition* operator on components. The product of two components consists in composing their corresponding layers separately. Parameters are used to define new interactions as well as new priority rules between the composed components. It allows *incremental* construction, that is, any compound component can be obtained by successive composition of its constituents. This is a generalization of the associativity/commutativity property for composition operators.
- It provides a powerful mechanism for structuring interactions involving both strong synchronization (rendez-vous) or weak synchronization (broadcast). Synchronous execution is characterized as a combination of properties of the three layers. Finally, timed components can be obtained from untimed components by applying a structure preserving transformation of the three layers.
- It allows considering the *system construction process* as a sequence of transformations in a three-dimensional space: *Behaviour* X *Interaction* X *Priority*. A transformation is the result of the superposition of elementary transformations for each dimension. This provides a basis for the study of property preserving transformations or transformations between subclasses of systems such as untimed/timed, asynchronous/synchronous and event-triggered/data-triggered.

The CEA, INRIA and Thales teams are contributing to the elaboration of a new standard: ***Executable UML foundation*** [MFJ05] that aims at providing a formal framework for defining an execution semantics of UML profiles in order to help harmonizing other standards. Its objective is to enable a chain of tools that support the construction, verification, translation, and execution of computationally complete executable models.

PARADES has been instrumental in transferring the knowledge of the Metropolis framework and related design methodology to a set of industrial designs and to the HRC modeling effort in SPEEDS. During the design of the industrial projects for PARADES partners (ST and United Technology), it was evident that the user-interface and architecture of Metropolis was intended for experts in the methodology supported by Metropolis and in the semantics of the tool. PARADES was instrumental in inspiring the transition form Metropolis to Metropolis II, where the architecture of the environment is intended to facilitate the job of the system architects and developers. PARADES' industrial nature was essential in this step. The principles upon which Metropolis II rests are mathematically the same as Metropolis but the implementation of the semantics is essentially different. In particular, the essential characteristics to be retained were:

- languages or conform to different models of computation.
- The capability of taking different parts of a design and refining/abstracting them such that these relationships can be verified.
- Relating together the architectural platform and the functionality in different ways to explore different realizations of the system. This design space exploration process may be carried out in terms of different metrics, such as throughput, latency, jitter, power consumption etc.

Like Metropolis the semantics of the Metropolis II framework will be centered around the connection and coordination of components. Unlike Metropolis, the components will be specified using external languages and the framework will serve to integrate these languages and their supporting tools. We use the same definitions for events, actions, and services as Metropolis. An *action* is a primitive concept. It roughly corresponds to a piece of code in the design. *Variables* (state) may be explicitly associated with an action. An *event* represents the execution of the beginning or the end of an action by a particular process. A *service* is a set of sequences of actions, with a unique begin/end event pair. Variables in the scope of the begin event can be used as service arguments. Variables in the scope of the end event can be used as return values. Events, and by extension, services, may be annotated by quantities of interest. Quantities capture the cost of carrying out particular operations and are implemented using quantity managers. *Quantity managers* are special components that provide annotation services. Schedulers are similar to quantity managers, but instead of a quantity they provide scheduling and arbitration of shared resources. Depending on the MoC used and the needs of the design, different quantity managers and schedulers can be used.

In Metropolis II designs are specified by instantiating and connecting different components, and then annotating and constraining their interactions. Metropolis II as Metropolis can describe with these primitive concepts both functionality and architectures. The role of quantity managers is essential in defining and manipulating non functional quantities. The links between functions and architectures needed to support their implementation is provided by the *mapping* mechanism that associates events between functional and architecture net-lists. Metropolis II is also intended to support mixed operational-denotational specifications. Constraints are expressed in the system using first order temporal logic and regular expressions. The execution semantics in Metropolis is provided by intersection of behaviors

and constraints. A simulator is intended to operate based on the operational description "filtered" by the constraints. Metropolis II supports non deterministic systems.

Metropolis can then support rich components and provides verification and synthesis services. In the future, the role of the various tools listed above in a loosely integrated platform will be carefully considered.

2. Platform for the analysis of safety critical embedded systems

The works carried out for this platform are building on UML profiles, in particular MARTE and HRC. The main efforts this year concern back-end tool chains, starting from one of the envisaged semantic level formats and integrating validation and code generation tools. The work on the front-end tools, providing mappings from user level profiles to semantic level formalisms has only started for MARTE and will start within the next year for HRC.

Two important collaborative projects for the platform have started this year which will provide the main contributions on this platform:

- The French National project OpenEmbeDD (http://openembedd.inria.fr), which includes the ARTIST Partners CEA, France Telecom, INRIA, Thales, and VERIMAG, work will start on mappings from the user level formalisms SDL and the MARTE UML profile to the semantic framework of BIP, developed at VERIMAG and to INRIA's Kermeta model for further connection with validation tools.
- The SPEEDS project IP SPEEDS, with partners INRIA, OFFIS, PARADES, and VERIMAG, has started this year. The work involves the development of a system level UML/SysML compliant framework for heterogeneous components, which will benefit from MARTE; it will be connected via semantic level formats like BIP to the validation platforms IF, Metropolis and RUVE.

The INRIA team developed a tool chain using tools of several ARTIST teams (mainly IF, Kronos, Giotto, Kermeta). The chain aims at supporting a complete software design process for real-time components, from service specifications down to executable software components in Java or C. The component implementation process uses a two-step method: designers construct an abstract implementation using timed automata, which is checked against the specification using the IF and Kronos tools from *VERIMAG*. A model driven approach is then applied to build a platform dependent implementation. The concrete implementation is generated by model transformations using tools from *INRIA Triskell team*. The target architecture is the Giotto platform designed by the *EPFL* team. The tool chain covers the life cycle of timed components from the service specification in timed tree logic down to algorithms coded in Java and executed on a Giotto runtime. The tool chain implementation has been done by INRIA and is now completed. A future meeting of the RTC platform group will include demonstrations and technical discussions on further integration of this chain in the platform.

The BIP framework that will be extensively used in several projects has been implemented in a tool: the tool consists of a front-end for editing and parsing BIP and generating C++ code to be executed and analyzed on a backend platform consisting of an engine and the infrastructure for executing the generated C++ code. BIP has been entirely implemented in C++ on Linux and uses POSIX threads. The execution engine iteratively executes the following step. At a given state, it monitors the state of atomic components and finds all the enabled interactions by evaluating the guards on the connectors. Then, between the enabled interactions, priority rules are used to eliminate the ones with low priority. Amongst the maximal enabled interactions, it executes one and notifies the atomic components involved

in this interaction. The notified components continue their local computation independently and eventually reach new control states.

The current implementation is suited for the state space exploration-based analysis of systems but presently not for developing embedded operating systems kernels and low-level services. This will be tackled in the next period and a connection to the analysis tools of the IF tool-set is ongoing.

BIP has been and is being used for modelling several smaller case studies and some larger on systems in the context of performance oriented systems satisfying hard timing constraints, in the context of planning tasks of autonomous robots and in the context of modelling energy consumption in sensor networks.

BIP/THINK collaboration between FTRD and VERIMAG has started this year. The goal of this BIP/THINK joint effort is to get simultaneously the executable embedded code of an OS and its behavioural model for analysis and verification. Today a transformation exists from BIP (used for analysis) to THINK. This project is now financed in a project in the context of EMSOC.

The UPPAAL tool for verification of timed automata has been upgraded by the Uppsala team for being able to handle UML specifications. It has been integrated in the Eclipse platform and the UPPAAL modelling language has been extended with hierarchical state machines, to support modelling of hierarchical structures and abstract behaviours of components. Presently, we are in progress to extend the UPPAAL modelling language with asynchronous communication channels. The idea is to use timed automata to describe the communication patterns and relative speeds of components in producing and consuming messages.

3. Platform for the analysis of performance critical systems

This platform is presently developed in the context of the French Persiform (http://www-persiform.imag.fr) project (with ARTIST partners FTRD, INRIA and VERIMAG). The aim of this project is the integration of performance evaluation and formal verification in requirement and design activities.

A first aim is to connect commercial performance analysis tool (event-based simulation mainly) to functional UML modelling tools for high-level performance analysis, in particular service specifications expressed in terms of activity diagrams [BCG*06] and sequence diagrams. For this purpose, a profile for the use of activity diagrams has been defined and a formal semantics has been through a mapping to a restricted class of coloured Petri nets plus annotations with probabilities and distribution concerning timing and resource usage. Annotated Petri nets are then transformed into performance evaluation platform SES Workbench (http://www.mmsolutions.com/english/workbench.htm). Alternatively MSC can be handled a transformation into the same class of annotated Petri nets. These transformations are based on the construction of meta-models for the different languages and transformation rules.

The initial chain has been applied to two industrial case studies on which the tool chain can be demonstrated.

For real time systems, work is performed on design specifications with performance annotations, compatible with the MARTE profile. These specifications will be transformed into IF models for validation of real time properties [HAB*05].

In the next period, mappings to tools for functional validation, in particular to IF or BIP are planned. As well as some work on representing observed traces during simulation by MSC. In the future, we consider mappings to other performance models.

4. Platform for the certification of smart-card applications

The work on this platform is supported by a collaboration between CEA and VERIMAG on functional validation of critical applications on smart cards. This work is carried out in the context of a national project, EDEN 2, that pursue the work done in the previous one, EDEN, in order to reach a consolidated implementation for industrial exploitation. It has not progressed exactly according to the plans which foresaw the definition of a UML profile for security properties, but it turned out to be more important to concentrate in the first year rather on the validation engine.

5. 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; this topic is covered by the Verification cluster and platform activities and the background tools of the platform have been described in the year 1 deliverable. We describe here on some new developments directly linked to the connection of existing verification tools to the modelling languages considered in the platform.

An interesting technical challenge for adapting **UPPAAL for asynchronous models** is to check the boundedness of channels, and to synthesize the maximal size of memory blocks needed to implement the channels. Preliminary results are reported in [KY06] showing that the expressive power of such systems with two channels -- that are no more expressive than finite-state machines in the untimed setting -- is Turing-equivalent. We have been developing methods based on approximations. As an abstraction for communication interfaces, we have adopted arrival curves from network calculus. Some preliminary results are achieved, and they will be implemented in the coming versions of UPPAAL for verification of systems with asynchronous communication. The work will be extended and integrated in the TIMES tool [FMPY06] for approximate schedulability analysis of systems with multi-resource and heterogeneous components.

The **symbolic execution kernel**, Agatha [GLRT06], has been extended to support analysis of heterogeneous model using **heterogeneous models of computing**. Developed by CEA through three national projects (STACS, Usine Logicielle and EDEN 2), it is implemented as an Eclipse component for test generation from UML models and within EDEN 2 project it is connected to the VERIMAG IF tool in the context of the platform 3 for certification of smart-card applications.

We developed sufficient criteria for guaranteeing properties of component systems by exploiting the structure of the BIP framework that strictly separates the description of behaviour of components from the way they interact and execute. We have considered so far liveness, local progress, local and global deadlock, and robustness [GGMSM06]. The criteria depend on different degrees of abstractions of the behaviours of the individual components and on the global interaction and priority model. We also investigated the incremental construction of proofs of such properties.

3.1.3 Objectives and Work Planned: Sept 2006 – February 2008

Globally, the work will continue according to the last 18 month plan and the tool chains which started to be developed will be further extended and/or connected. The initially planned tool integration through the jETI tools is likely not to happen within the next 18 months, although an interesting future perspective. Also the work on the platform for the certification of smart-card applications is likely to be less important than initially foreseen.

In the next period, the work will concentrate on enlarging the existing tool chain kernels by means of new model transformations, and by bringing the modelling standards closer together.

Platform for the analysis of safety critical embedded systems

The main future work on this platform will be carried out within the projects System@tic/Usine Logicielle, OpenEmBeDD, ATTEST and SPEEDS. They will concern the missing connections between the modelling languages used and the back-end tools via semantic level intermediate formats. It concerns also some work on back-end tools which are specific for this platform.

UML analysis tools will be extended to support the HRC (heterogeneous rich component) models developed in SPEEDS. The HRC model will be mapped to semantic level formalisms so as to allow analysis and validation of such models by existing tools, in particular those developed by INRIA, OFFIS, Parades and VERIMAG. The analysis techniques will comprise *compatibility checks* on composition of HRC design units, including static checks as well as verification of connections of assumption/promise pairs as well as *refinement verification* using the verification of black box specifications against grey box specifications. Furthermore, one has to check whether the black box specification combined with all assumptions imply the promises of the component. Based on the rich component approach the analysis techniques will extend the aspects of systems covered from behaviour and real-time to encompass also safety and other non functional aspects, in particular those important for supervising the system development process.

The Kermeta-IF-Giotto prototype tool chain already *manages software development support starting from the specification of components and their composition, their verification down to the generation of Java or C based executable units for the execution of a specific platform; it* will be disseminated, strengthened and improved thanks to the platform participant's feedback. Furthermore, an integration work of the Kermeta-IF-Giotto chain with the SPEEDS semantics (see below) will start in November, 2006. It is expected that the chain will be merged into the larger SPEEDS platform in 2007.

The Metropolis II framework will be developed by PARADES in collaboration with several external partners such as the University of California at Berkeley, United Technology, Cadence and ST who will provide important test cases for system level design crossing company boundaries. In particular, design space exploration with multiple complex architectures described with functional and non functional properties will be addressed. Industrial applications will include automotive, wireless sensor networks, industrial control and multi-core chips.

The BIP engine will be connected with the analysis tools of the IF tool-set and analysis techniques using the particular structure of BIP specifications will be implemented and extended so as to provide a suitable verification and analysis engine for system level models provided by SPEEDS case studies. For the connection to the HRC meta-model, we consider reusing the Kermeta-IF tool chain.

The BIP/THINK tool chain represents a back-end of a tool chain for code generation for given platforms. This chain is intended to be used in combination with the validation backends for BIP and for models imported through front-end tools. Code generation is not directly in th objectives of SPEEDS, but the existence of some code generators may turn out to be very useful for demonstration purposes. Future work concerns both improving the existing compilation chain: some existing OS for small targets and sensor nodes (TinyOS, SOS, Nano-RK, Mantis ...) will be considered to check how they behaviour can be modelled using BIP formalism. Simultaneously, a component-based architecture of these systems will be proposed using the THINK Component Based Framework. This architecture will be completed by THINK control components obtained from the previously mentioned OS BIP model using the BIP/THINK translator developed this year. Finally, the resulting architecture will be implemented on concrete hardware platform using the available THINK environment.

The ultimate goal is making available these validation techniques, as well as code generation techniques to the designers in commercial tools, in particular those considered in SPEEDS, that is SCADE and Rhapsody. We expect that the work done within the SPEEDS project will contribute to a stronger integration of tools.

The implementation of development and validation frameworks building on the MARTE profile include also model transformation and the code generation of a concrete implementation of the Accord/UML modelling and design platform developed at CEA. Cantabria will study the appropriate level of abstraction to extract transactional analysis models. These models will then be used to apply the schedulability analysis and performance evaluation tools that are developed inside the MAST suite. Along the process it may result necessary to adapt, extend or restrict the applicable capabilities in any of the two modelling platforms, so both are subject to adaptation in the search for complementarities. The work for the next months will include revising and if necessary proposing methodological/practical strategies to the use of concrete components technologies. The first to be considered will be an implementation of RT-CORBA and then the combination of the Real-time and Distributed annexes of the new Ada2005 programming language, which includes now the capability to describe interfaces. Based on the achievement, we will further develop validation techniques taking into account the characteristics of the modelling languages used in the platforms, in particular for properties related to the interaction in component-based systems. In the OpenEmBeDD project, the connection with several back-end tools is planned. Some of those mentioned in the context of SPEEDS, will be made available in OpenEmBeDD. Joint use of both profiles is envisioned at a later time.

Integration between MARTE standard and automotive domain, and in particular with Autosar standard, will be continued in ATESST project. The EAST-ADL profile will be extended during next year in order to ease the modelling of product families (or product lines) by adding elements of variability description, in particular for variation of component behaviour.

Another line of work on the MARTE profile will be on its integration along the whole system development process through defining **traceability support for UML** based development in embedded system. Based the three UML profiles SysML, MARTE and EAST-ADL 2, an Eclipse component will be developed within the MemVaTEx French project. The project is strongly coordinated with ATESST by its leader, Siemens VDO, and involved in particular CEA and INRIA cluster partners (see www.memvatex.org).

Platform for the analysis of performance critical embedded systems

An important part of the work in the next period until the end of the Persiform project will consist in consolidating the existing tool chain and in evaluating it on hand of more extended case studies.

In addition, mappings to tools for functional validation, in particular to IF or BIP are planned. Such mappings may also be used for representing observed traces during simulation by MSC. It is also planned to set up a follow-up project, in which mappings to other performance models and a stronger integration with the design process is envisioned.

Platform for the certification of smart-card applications

The work on this platform continues to be carried out in the EDEN-2 project and will mainly port on functional validation of critical applications on smart cards. The definition of a UML profile for security properties is considered for the third year of the project.

Collaboration and Dissemination

Like already this year, we will privilege open meetings and organisations of workshops over cluster meetings in a close format. We are again organising the MARTES workshop with MoDELS 2006 in Genoa. As the workshop attracts an increasing number of submissions and participants (this year 40 participants are expected), it will probably be organised again in 2007. It is also planned to organise a platform workshop as a satellite workshop of an appropriate major conference.

3.2 Cluster Integration: Development of UML for Real-time Embedded Systems

Activity leader: Sebastien Gérard (CEA)

3.2.1 Year 1 Achievements: Sept 2004 – August 2005

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

- The first objective of this first period was first to 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 in the context of the Real-time, Embedded, and Specialized Systems (RTESS) Platform Task Force in February 2005: http://www.omg.org/cgi-bin/doc?realtime/05-02-06 .
- The second objective was 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 provider) involved in this aspect at the OMG. It is composed of: Artisan, Carlton university, CEA, IBM, I-Logix, INRIA, Looked-Martin, Thales, Tri-Pacific. We continue this effort in order to improve the force of our consortium.
- Finally, a framework for the unification of the two analysis sub-profiles in the original SPT profile was proposed and some effort was made to simplify the way final models will be annotated.

In the context of the Omega project, a UML profile has been developed appropriately 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. A successful profile has been defined in [GOO05] and successfully applied in the Omega project [OME05, OGOL05a, OGY05] and elsewhere [HCBA05]. This work yields maximal expressiveness and semantic level foundation.

Finally, during this first year period, we tried as much as possible to give all the required information about OMG standardisation processes and specially points related to RT/E to all ARTIST partners.

3.2.2 Year 2 Achievements: Sept 2005 – August 2006

A consolidated architecture for the Marte profile
The Marte profile architecture model consists of three main packages:

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

– The RealTime and Embedded application Modeling package (RTEAM); it enables modeling of RT/E application. It concerns mainly defining in one hand high-level model constructs to depict real-time and embedded features of application, and in other hand to enable the description of execution platforms, software as well as hardware.

– The RealTime 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.

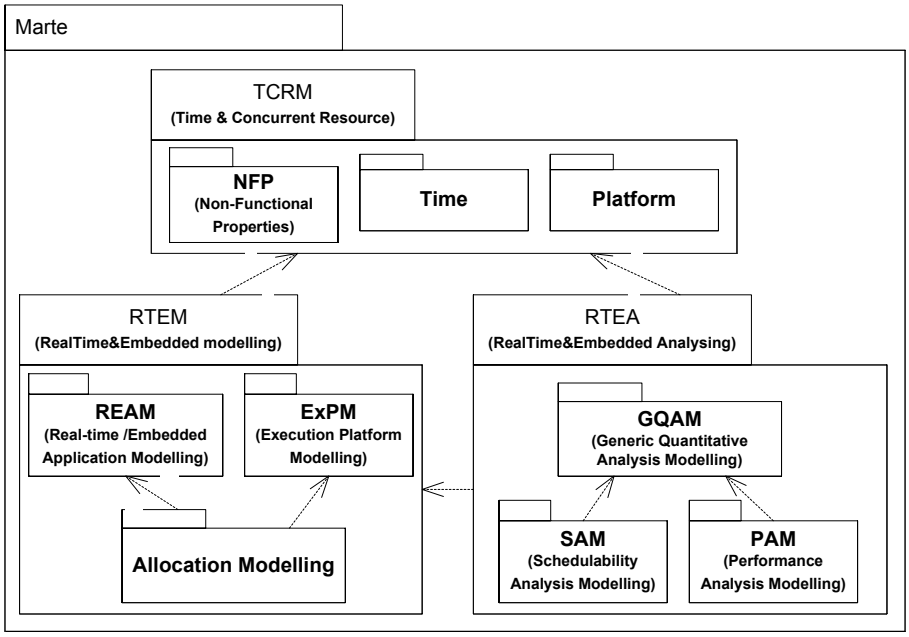

Figure 1. Current architecture of the Marte profile.

3.2.3 Objectives and Work Planned: Sept 2006 – February 2008

Before explaining the job scheduled for this activity within the next 18 monthes, let's remain what is the OMG standardization process.It consists mainly in four stages:

- Request For Proposal (RFP): this satge consists in defining a set of requirements that will be considered has to be considered in a new standard. For Marte, it is the document referenced as: realtime/05-02-06.
- Initial Submission (IS): this document is time to OMG consortium to post their intend to answer to the RFP. The level of granularity of this document is not defined. It can be very detailed but also only a high level description of the work intended to be provided by the consortium. In the case of marte, only one consortium has declare its intend to provide a solution to the Marte RTF, the proMarte consortium (www.promarte.org). CEA (chairman of this consortium, which is also co-chaired by Ben Wtason from Looked Martin), Thales and INRIA are part of this consortium and are main of the contributor through the CARROLL-PROTES project (http://www.carroll-research.org/uk/projets/projets.htm). The Marte IS is referenced as realtime/05-11-01.

- Revised Version. This is a detailed document considered as the future OMG standard for a given RFP. This document is part of the job planned in the next 18 monthes of work of this activity.
- Final Version. Once the revised version has been voted, an OMG group (usually build from the consprtium that has proposed the revised version of the standard) is build in order to manage the finalization task force for the standard. Within this period, the standard is made available to other people than these one of the consortium, and these reviewers can raise issues against this standard in order to debug/improve this latter. At the end of this period, the FTF produce the final version (forMarte, the v1.0). This period is only dedicated to debug the standard and is not place for introducing new concepts! For that purpose, people will have to wait the next iteration in the standardization process called the revised task force that will produce the version 1.1.

So, in this context, the next 18 monthes period will be split into 2 parts:

- Until March 2007, we will continue on the writing of the standard itself. The vote for this standard is scheduled to happen Q1 of 2006. In parallel, we will have to make some proof of concepts defined in this standard. The version of the document that should be provided in March 2007 is called a revised version (the version after the initial submission).
- From April 2007 untill Feb 2008, we will work on applying this standard and give feedback to OMG in order to prepapre the final version that should be due by mid of 2008. In addition of that, we will spent a lot of effort for the dissemination of the standard in order this latter may be considered by industrials and academics in their future projects and researches.

3.3 Cluster Integration: Component-Based Design of Heterogeneous Systems

Activity leader: Bengt Jonsson (Uppsala)

Following the decisions at the Review in November 2006, this new activity will subsume the research activities and meetings that in Year 2 was found in the activities

- *NoE integration: Forums with specific industrial sectors.*
- *NoE integration: Seeding new research directions.*

This activity is structured into three sub-activities:

- Design of Heterogeneous Systems (leader: Joseph Sifakis)
- Interfaces and Composability (leader: Bengt Jonsson)
- Industrial Liaison (leader: Werner Damm)

3.3.1 Year 1 Achievements: Sept 2004 – August 2005

This activity started at the end of Year 2.

3.3.2 Year 2 Achievements: Sept 2005 – August 2006

A detailed description of these achievements is provided in the RTC Cluster deliverable and in the activity deliverables of Forums with specific industrial sectors *and* Seeding new research directions.

Design of Heterogeneous Systems

The theory on tag systems has been further developed by Benoît Caillaud and Dumitru Potop-Butucaru (VERIMAG, then INRIA, team Aoste), who have developed a theory for the correct deployment of synchronous designs over globally asynchronous, locally synchronous (GALS) architectures. This work introduces 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 is introduced. The model covers classical implementations, where a notion of global synchronization is 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 researches on 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 has been 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. Work on the integration of existing validation techniques, implemented in the IF platform, is ongoing. 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 SAT checker modulo scheduling theory 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, who are currently 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* is being developed, allowing for cross-viewpoint assume-guarantee reasoning.

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. Techniques from schedulability analysis are further developed by *Cantabria* and *Thales* in the newly started project FRESCOR: 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 has been developed]. 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

Here we summarize the most important conclusions from this meeting.

Regarding the interaction *contro/embedded software*:

- There is a permanent misunderstanding between control & software engineers
- Regarding the relative merits of ET/TT, control design aspects provide complementary views, not considered before
- There is a need for a notion of component for control that would enable incremental development of control systems.

Regarding AUTOSAR:

- The AUTOSAR design flow for distributed embedded electronics is not completely plug-and-play, neither it is compositional, for reasons of scheduling: scheduling is, today, based on global systems models. Component-based techniques for real-time are needed. (This is an ongoing research activity at some ARTIST2 teams participating to RTC and Executions Platforms clusters.)
- Turning the AUTOSAR approach into an effective tool for dispatching the work efficiently among suppliers is still seen as a challenge.

3.3.3 Objectives and Work Planned: Sept 2006 – February 2008

Design of Heterogeneous Systems

Leader: Joseph Sifakis (Verimag)

The pursued work directions are mainly these identified in the Workshops organised in Seoul (http://www.artist-embedded.org/artist/Overview,29.html) and Zurich (http://www.artist-embedded.org/artist/MoCC-06.html) sulmmarized as follows:

- Investigate relations between software engineering “object-oriented” views for components and “system-oriented” views. The first consider components as a means for structuring functions and data. They support point-to-point interaction model mainly by function call. The second consider components with behaviour and rich interfaces. They support a variety of interaction and execution models e.g. Ptolemy.
- Study unifying semantic frameworks for “system-oriented” components. We distinguish two actions lines. One in the continuation of the work by INRIA, Parades and Verimag for the unification of models of computation based on denotational semantics (tagged traces). The other based on operational semantics in the continuation of work pursued mainly by Verimag (BIP).
- Study notions of expressiveness taking into account structure for the comparison of existing component frameworks.

The theoretical studies developed along these lines will be validated by work on methods and toolsl carried out in industrial projects such as SPEEDS.

Interfaces and Composability

Leader: Bengt Jonsson (Uppsala)

The pursued work directions are mainly these identified in the Workshops organised in Seoul (http://www.artist-embedded.org/artist/Overview,29.html) and Zurich (http://www.artist-embedded.org/artist/MoCC-06.html) sulmmarized as follows:

- Study theory and rules for correctness-by-construction. Two classes of rules will be inverstigated. The first includes compositionality rules for building correct systems for correct components. The second includes composability rules ensuring that essential properties of components are preserved along integration.
- Study concepts of interfaces and associated (partial) composition operations. These operations should encompass interface compatibility relations where non compatibility means violation of simple behavioural properties e.g., deadlock-freedom.

- Investigate the application of Assume/Guarantee techniques to component frameworks. This work direction is strongly related to the previous one and focuses on compositionality of non functional properties.
- Investigate the application of abstraction techniques to component frameworks. This work direction is strongly related to the previous ones and focuses on the computation of abstractions for non functional properties.

The theoretical studies developed along these lines will be validated by work on methods and toolsl carried out in industrial projects such as SPEEDS.

Industrial Liaison

Leader: Werner Damm (Offis)

Further exploiting the already organized meetings

The conclusions of the meeting *Beyond AUTOSAR* will be organized in the form of a paper for submission to a scientific magazine with large audience in the targeted industrial sector. Preparation of this paper already started and is supposed to last for the rest of 2006, and possibly early 2007. This paper will be prepared by a group of academic participants to this meeting, who staid for the following Saturday in Innsbruck to prepare its contents. We hope that, despite the overloading of key authors of this paper from various sources and projects, the paper will be available for the next report.

Planning additional meetings

A subsquent meeting of this activity is planed to be held on the aeronautics/avionics sector, in summer/autumn 2007. The preparation and exploitation of this latter meeting will allow for an analysis of the situation of and challenges raised by Integrated Modular Avionics (IMA).

4. Cluster: Adaptive Real Time

Cluster Leader: Giorgio Buttazzo (Pisa)

The following is a description of the activities and overall objectives for the period: September 2006 – February 2008. The next reporting period will cover September 2006 – August 2007.

Staff Mobility

None reported.

4.1 Platform: A Common Infrastructure for Adaptive Real-time Systems

Activity leader: Giorgio Buttazzo (Pisa)

4.1.1 Achievements: Sept 2004 – August 2005

After analyzing the state of the art of real-time kernel technology, the Shark operating system (http://shark.sssup.it/) was selected to be used as a shared platform in the ART cluster for experimenting novel real-time algorithms for control applications. Shark was selected because its modular structure enables the user to easily replace the scheduler or the mutual exclusion protocol with a different one, without changing the application code. Moreover, it supports applications with explicit timing constraints, it includes several advanced algorithms for task scheduling and shared resource management, which can be dynamically selected by the user through a configuration file, it includes drivers for the most common I/O peripherals, it complies with the POSIX standard (PSE51 profile) and includes user manuals and several sample real-time applications. Finally, Shark was developed at the ReTiS Lab of the Scuola Superiore Sant'Anna of Pisa, in collaboration with the Robotic Laboratory of the University of Pavia, hence the know-how for maintaining and updating the kernel is internal to the ART cluster.

The University of Pavia, in collaboration with the affiliates Evidence and the Scuola Superiore Sant'Anna, developed a web site (http://feanor.sssup.it/retis-projects) to create a forum for the various Shark users, in which it is possible to exchange messages, search for questions, etc. The web site also includes a page with web links to the various research groups that are using the kernel for control applications, and a page of *Frequently Asked Questions*, to quickly address the most common problems encountered by the developers.

The main activity was to organize a workshop to introduce the Shark kernel to all the partners of the ART cluster, enabling the participants to quickly use the kernel, develop simple real-time applications, and implement novel scheduling algorithms. The workshop was held in Pontedera, Pisa (Italy), at the Scuola Superiore Sant'Anna, from February 28 to March 4, 2005. During the workshop, participants were asked to develop a small software project using Shark. They were divided into groups and each group had to develop a software module or at least design the structure of the application. The time available for programming was not much, therefore some group developed some simple demo, while some other participant just sketched the idea of a more complete kernel component.

4.1.2 Year 2 Achievements: Sept 2005 – August 2006

A detailed description of these achievements is provided in the activity's deliverable.

A new kernel release: Shark 1.5.1
A new kernel release, Shark 1.5.1, was issued on July 25, 2006, introducing several new features with respect to the previous version of the kernel (Shark 1.5). The work has been carried out at the University of Pavia, in collaboration with the Scuola Superiore Sant'Anna and Evidence S.r.l. The new kernel includes the following additions:

- Event filters have been added to the Tracer to allow the user to select the events to be monitored at runtime. Moreover, the tracer output can now be saved on local disks, in addition to remote servers.
- The support for the USB has been introduced. The support for Host and Hub devices is fully working; USB mouse, keyboards, joystick and joypad are already fully supported, while PWC chipset-based webcams are working with negligible problems on some specific machines.
- The IntDrive interrupt server for Linux-imported drivers has been updated in order to correctly implement the ideas proposed in the related scientific paper. There is now a new interface for initialization, and some delicate internal mechanisms has been fixed.
- A support for Dynamic Voltage Scaling has been provided to allow S.Ha.R.K. to exploit the power management techniques available in the most recent processors. AMD PowerNow and Intel Centrino SpeedStep are currently fully supported. The feature has been implemented as a kernel module and it is fully compliant with the Linux CPUFreq driver.
- A S.Ha.R.K. Quick Guide has been released to simplify the work to the new users. It covers the basic topics for getting familiar with S.Ha.R.K., such as installation, basic application development, system configuration, and remote execution of S.Ha.R.K. applications to improve productivity.
- Other new documents include the manual for the new S.Ha.R.K. Tracer (which covers the initialization, usage, and log saving procedures) and the HTML and TXT versions of the kernel manual. In particular, the TXT version is useful while developing under DOS, where both PDF and HTML documents cannot be handled.

URL: http://shark.sssup.it/
URL: http://shark.sssup.it/repository/shark

A repository of real-time applications
A repository for all real-time applications developed using the Shark kernel was created to facilitate the users in the development of new real-time software. The repository includes a a folder of supported applications and a folder of all unsupported software. The supported applications are tested and maintained by the developers to be consistent with the current kernel version, while the unsupported folder includes all programs, demos, and advanced applications developed under Shark and made available by the maintenance team.
URL: http://shark.sssup.it/repository/applications
URL: http://shark.sssup.it/repository/unsupported

A repository for scheduling modules and resource management
A repository of all kernel modules developed for the Shark operating system was created to facilitate the users in the development of new kernel mechanisms. The modules include scheduling modules (implementing periodic schedulers, aperiodic servers, or overload management policies) and resource modules (implementing concurrency control protocols for accessing shared resources). Each modules contains the C source code and required headers compliant with the Shark module specifications.
URL: http://shark.sssup.it/repository/modules

Shark at University of York
This work has been done by Alex Zerzelidis at the Computer Science Department of University of York. The Preemption Level Protocol (PLP) allows EDF to be used on top of priority queues and was introduced in

> Burns, A., Wellings, A.J., and Taft, T. S., "Supporting Deadlines and EDF Scheduling in Ada", Lecture Notes in Computer Science, Springer-Verlag, Volume 3063 / 2004.

Alex Zerzelidis extended the PLP to include multi-unit resources. To achieve that, two amendments were introduced to the original protocol:

Amendment 1 (applies to point 2 of Section 3) Resource ceiling preemption levels *are dynamic* and any resource *R* has a current ceiling $\lceil R \rceil$ defined as

$$\lceil R \rceil = \max(\{0\} \cup \{\pi(\tau) \mid \nu_R < \mu_R(\tau)\}) .$$

where $\pi(\tau)$ is the preemption level of task τ, ν_R denotes the number of units of R that are currently available and μ_R is the maximum requirement of task τ for R. That is, the current ceiling of resource R is the maximum of zero and the preemption levels of all the tasks that may be blocked directly when there are ν_R units of R available.

Amendment 2 (applies to point 4 of Section 3) When a task accesses a resource, the current ceiling level of the resource is recalculated and the task's active priority is raised to the new current ceiling level of the resource. The active priority of all other tasks already accessing the resource remains unchanged. When the task releases the resource the ceiling is recalculated.

The PLP and its multi-unit extension was implemented on the SHARK RTOS as a new scheduling module in the kernel, called MPLP (Multi-unit PLP), which is based on the EDF and POSIX modules. The module can be registered as usual in the function

```
TIME __kernel_register_levels__(void *arg)
```

if we include a call to `MPLP_register_level(int flags, int prioritylevels)`.

The module implements PLP on a set of priority levels (their number specified by the `prioritylevels` parameter) and allows tasks to share multi-unit resources. Tasks can be created at runtime and declare a list of resources to use.

Shark at the Technical University of Kaiserlautern
Shark has been used at the Technical University of Kaiserlautern (TUKL) as a platform for video processing applications and, in particular, for experimenting flexible resource management policies for user quality control. The methods estimate resource demands for frame decoding and use Shark scheduling algorithms for guarantees.

At TUKL, the Shark kernel was also adopted for education to teach real-time scheduling in an undergraduate course. Two labs are being developed: in one lab, students are required to develop a scheduling algorithm, analyse it and test it on the Shark kernel; in the other lab, students have to implement a simple video processing algorithm running on Shark, learning implementation and overhead issues.
URL: http://rts.eit.uni-kl.de/research/adaptive-rts

Shark for control applcations
A contact has been established with Dr. Sjur Vestli, from Logobject AG – Switzerland (http://www.logobject.ch/), for a possible use of Shark in robotic applications, and in particular for the control of autonomous vehicles. The group is currently using a fixed priority kernel for Power PC platforms, but they are moving to PC architectures and are looking for a kernel supporting dynamic scheduling and advanced resource management techniques for an efficient exploitation of the onboard resources.
URL: http://www.logobject.ch/

Shark for education in a Master Course for IIT graduate students
Shark was used as a sample real-time kernel in a Master course on Real-Time Systems organized in Pisa by the Scuola Superiore Sant'Anna, from March 2006 to May 2006, for 20 Indian graduate students selected from the India Institute of Technology (IIT). Students worked in groups of two or three people to develop a number of real-time concurrent applications and make experience in using advanced scheduling techniques.

Shark for education at University of Catania
Shark was used as a sample real-time kernel in a graduate course on Real-Time Systems given by Lucia Lo Bello at the University of Catania, from March 2006 to June 2006. Students worked in groups of two or three people to develop a number of real-time concurrent applications and make experience in using advanced scheduling techniques. A description of the projects is available at
URL: http://www.diit.unict.it/users/llobello/retisnetlab/shark.htm

Shark at the University of Illinois – Urbana Champaign
A testbed was developed at the Real-Time Systems Lab of the Computer Science Department of University of Illinois at Urbana-Champaign to show the applicability of real-time OS Shark and real-time MAC protocol RI-EDF for distributed control applications, where sensors/actuators are connected through a wireless channel. A wireless distributed control system for an inverted pendulum was built as a testbed environment. In the application, camera sensors are used in complement with potentiometer sensors on the cart to balance the pendulum pole. Real-time sensory acquisition was performed on a PC using the S.Ha.R.K. operating system <http://shark.sssup.it/>, whereas Berkeley Mica2 motes running TinyOS <http://www.tinyos.net/> were used for wireless communication with the RI-EDF protocol. We were able to successfully control the pendulum and utilize the extra network bandwidth for other real-time communications on the same shared channel using RI-EDF protocol. More details can be found at
URL: http://pertsserver.cs.uiuc.edu/~mcaccamo/IPC/index.html

Shark for robot control at University of Dresden
A contact has been established during ECRTS 2006 with the robotic group at the University of Dresden, for a possible use of Shark in robot control applications. The group is currently using a simple priority-based kernel developed in their lab for Motorola 68020 architectures, and then adapted to run on Power PC platforms.

4.1.3 Objectives and Work Planned: Sept 2006 – February 2008

In the next 18 months we are going to use Shark for understanding how to build a component-based operating system, where most of the available kernel features can be composed together to create several user-defined configurations.

In particular, a component based approach should separate mechanisms from policies in order to replace a scheduling algorithm or a resource management protocol without affecting the applications and the others components. In addition, it should allow combining different scheduling disciplines to support the development of hierarchical software architectures.

There are several benefits in adopting a component based approach at the operating system level. First of all, it would be possible to enhance the functionality of the kernel by adding new blocks, depending of the application requirements, so tailoring the kernel to the specific system to be developed. Secondly, it would facilitate and speed up the integration of novel research results, which could increase efficiency and/or predictability. Finally, it would simplify the process of porting the kernel on different platforms, so reducing the time to market and the development costs on upgrades (since only small parts should be developed).

However, there are several practical and theoretical problems to be solved, since most of the mechanisms implemented in a kernel (like scheduling, resource protocols, interrupt handling, aperiodic servers, synchronization and communication) heavily interact with each other and have a high degree of inter-dependencies.

We plan to treat such problems by addressing the following issues:

- decoupling scheduling algorithms from applications;
- decoupling scheduling mechanisms from scheduling policies;
- decoupling scheduling algorithms from resource management protocols;
- combining resource reservation with resource management protocols.

4.1.4 Meetings Planned

Meetings for Year 3 will be decided at the annual consortium meeting, Nov 7th in Paris (just before the review).

4.2 Cluster Integration: Flexible Resource Management for Consumer Electronics

4.2.1 Short Description

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, an experimental framework needs to be built in order to measure the performance of the different scheduling algorithms, and evaluate their applicability to real application domains

4.2.2 Year 1 Achievements: Sept 2004 – August 2005

From merged "Adaptive Resource Management for Consumer Electronics" activity.

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.

From merged "Flexible Scheduling Technologies" activity.

The work on flexible scheduling has focused on building a flexible scheduling framework that can be used to safely mix real-time and non real-time processes, and manage the available resources using the necessary flexibility to support the desired quality of level and maximizing the resource usage. Management of different resources has been studied in the previous work. New methodologies for integrating overload management techniques with energy-aware strategies have been developed. Resource reservation mechanisms that reduce intertask interference and provide temporal protection among the concurrent activities have been proposed. Techniques for integrating offline and online scheduling have been developed, and in particular on the coexistence of fixed priority and table-driven scheduling. The flexible scheduling techniques have been adapted to distributed systems by providing the ability to make dynamic bandwidth reservations using a distributed acceptance test that ensures the overall network schedulability.

Work has also been performed to ensure that the techniques developed can be used in different application domains. Work has been done towards using flexible scheduling techniques in adaptive resource management for media processing. Image processing applications associated to robotic vision and control, has been studied, a situation that is known for its varying load; in particular work for this environment has focused on the online adaptation of the processes attributes to enhance the quality-of-service of the application. A dynamic adaptation scheme has been developed for the control part of distributed embedded systems, which reduces the bandwidth requirements while maintaining the control performance. Implementation techniques for the developed novel scheduling techniques have been designed to port them to kernels that have to run on small microprocessors with scarce resources.

4.2.3 Year 2 Achievements: Sept 2005 – August 2006

From merged "Adaptive Resource Management for Consumer Electronics" activity.

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 will feed 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 have continued work on integrating real-time scheduling and cache management on multiprocessor platforms. To this end, we have carried out experiments to study cache behaviour on the actual platform and formulated a number of scenarios with increasing complexity. The work is being carried out by a joint PhD student.

Server Based Flexible Scheduling
Schedulability analysis techniques 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 has been focused on hierarchical scheduling analysis and design techniques A further issue has been the dimensioning the parameters of a server for minimizing the average response time of the served activities. A statistical approach has also been 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
http://rts.eit.uni-kl.de/research/adaptive-rts
http://cordis.europa.eu/fetch?CALLER=PROJ_IST&ACTION=D&RCN=62751&DOC=1&CAT=PROJ&QUERY=1158229719107

Adaptive resource management for networks
Work concerned the analysis of the achievable QoS guarantees in wireless networks. After defining a proper model for the main resources (i.e., CPUs, disk and network), a number of existing scheduling algorithms for the three types of resources have been analysed under fluctuating workload to evaluate their behaviour in terms of service guarantee. Then, the achievable end-to-end QoS guarantees have been 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, which is still an important networking technology in the field of distributed multimedia systems. A wireless 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 have 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
Complex dynamic real-time scenarios may prevent the possibility of computing optimal service configurations before execution, an iterative refinement approach with the ability to trade off deliberation time for the quality of the solution was specified. The approach is to quickly find a good initial solution and to propose heuristic evaluation functions that optimize the rate at which the quality of the current solution improves as the algorithms have more time to run.

The work has 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 is used to dynamically adapt future stability periods, according to a feedback control scheme. Work was done by Madrid and Proto (www.hurray.isep.ipp.pt/activities/qos)

Middleware
System adaptation requires full knowledge of the system state, therefore work has also been carried out in a framework to gather actual resource usage information, and interact with the operating system, extending the traditional POSIX trace model with a partial reflective model for operating system monitoring. The work was done by Porto.

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 higher layers are meant to deal with quality, while lower layers are mainly related with resource management. In particular, the lowest layer is intended to manage budgets or resource shares assigned to applications. This layer has to provided accounting and enforcing facilities to ensure that budgets are guaranteed. Some times this functionality is provided by what is called resource kernels.

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 is under development is 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

From merged "Flexible Scheduling Technologies" activity.

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. See a description in the section on workshops, below. The workshop was very sucessful and brought together 20 participants from the following institutions:

- Technical University of Kaiserslautern, Germany
- Visual Tools, Spain
- University of York, UK
- Thales Communications France
- Evidence, Italy
- University of Cantabria, Spain
- Scuola Superiore Santa Anna Pisa, Italy
- Technical University of Madrid, Spain
- Polytechnic Institute of Porto, Portugal
- CEA, France
- Rapita Systems, UK
- University of Aveiro, Portugal
- Czech Technical University in Prague, Czcech Republic
- Technical University of Valencia, Spain

The results of the workshop, a wide collection of application requirements, are now in the process of being refined to produce a final document with a clear set of requirements for the integrated resource scheduling framework.
http://www.artist-embedded.org/FP6/ARTIST2Events/Events/flex_sched/

Baseline for integrated-resource scheduling framework
The FIRST IST project that finished in 2005 produced as its main result a contract-based scheduling framework 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 has been selected as the baseline for the more ambitious framework that will be developed in this activity and that will take into account the integrated scheduling of multiple resources. The FTT framework 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. Moreover, the impact of flexible scheduling on dependability for distributed safety-critical applications was also assessed using FTT-CAN and appropriate mechanisms have been developed.

New theoretical developments
The contract-based scheduling framework defined above 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 has been developed towards being able to analize such scheduling schemes. In particular, methods have been developed in the University of York to analyze threads scheduled by fixed priority high-level schedulers based on different kinds of underlying fixed-priority schedulers. These methods have been extended by the University of Cantabria to EDF high-level schedulers based on constant-bandwidth servers. As a result, a complete analysis method exists for this kind of hierarchical scheduling.

Work has been 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. CAN is now widely used is a range of real-time systems including high-integrity and flexible and adaptive systems. Unfortunately the analysis usually applied to its use has a fault that means that in some circumstances a system can be deemed schedulable when in fact it is not. This problem can arise because of the non-preemptive nature of the CAN protocol, and manifests itself when the first release of a message at a critical instant is not the worst case.

SSSA has 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 has 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. These results will be published in year3

Flexible architectures and communication protocols for networks used in distributed embedded real-time systems
This research has been 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. This activity addressed the issue of achieving flexibility in networked subsystem integration through an integrated approach, where several systems are encapsulated as subsystems and later integrated on a shared hardware architecture. As the network is a resource shared by all subsystems in the distributed architecture, its role in the integration process is particularly important, and the usage of an efficient and flexible network scheduler is essential. The core activity thus focused on assessing the suitability of Server-CAN, a network scheduler for the Controller Area Network, in the context of subsystem integration.
- Facilitating subsystem integration by decoupling priority and identifier in CAN messages. The CAN message identifier is especially important, as it not only does identify the message, but it also determines the message's priority. Hence, special care needs to be taken when assigning identifiers to messages. The research done on this topic resulted in a paper which outlines how CAN based systems are engineered today and indicates the potential and benefits of decoupling the message priority from the message identifier. Three solutions to this are given: TT-CAN, FTT-CAN and Server-CAN, and their strengths and weaknesses in an integration context are discussed.

4.2.4 Objectives and Work Planned: Sept 2006 – February 2008

In the next 18 months we will expand the application domain of the activity to more general media processing, to provide for more industrial input, including non mass market video processing and telecommunication.

We will continue to collect requirements to feed input to development of our adaptive methods, including the expanded application domain.

We will expand the integration of resources to be managed jointly. With respect to scheduling and cache, we will develop first algorithms to reflect on both scheduling and cache management. We will develop algorithms for adaptation of fluctuating resources, in particular wireless bandwidth and stream transformations.

The integration of HOLA-QoS with SHARK and MARTE will continue during this period of time, by UC3M and UPM. In the context of this activity, the aims will be to test the overall environment with multimedia applications to check the fulfilment of the quality properties. The used adaptation algorithms will be revised in order to take advantage of the resource management features provided by SHARK and MARTE.

UPM and UC3M will continue their work on extending Java-RMI with QoS management features. The goal is to improve remote invocations predictability with respect to selected QoS properties and to ensure that a server has enough resources to attend to a given number of invocations from a set of clients.

Further investigation will be done on resource reservation mechanisms, in order to extend them to other types of resources (not only processors) and to make them work under mutual exclusion constraints and the possibility of using the elastic scheduling theory in a energy-aware context, with the objective of balancing energy consumption with bandwidth requirements.

Work on the use of kernels developed by partners for HOLA-QoS will continue and focus in particular on the integration of communication aspects. Design of a component intended to adapt budgets at the low level API of HOLA-QoS to the functionality and model of the budgets of SHARK and MARTE. Implementation of this component.

Energy management: optimization of power consumption for CE devices is desirable. Adjustable CPU speed and memory speed can provide mechanisms for prolonging batteries life. The adaptive tuning of both speeds will be analyzed. UPC will carry out this work.

The work on anytime, iterative algorithms for QoS provisioning in collaborative embedded systems, and will address server-based scheduling approaches for collaborative service isolation will continue.

4.2.5 Meetings Planned

Partners will meet regularly on the listed topics. In addition, the following meetings are planned.

Evaluation of the integration of flexible resource management techniques

Meeting Title	Evaluation of the a common flexible scheduling framework
Date	June 2007
Objectives and expected output	The main objective of this meeting is the evaluation of the framework for the integration of the management of all the system resources involved in an embedded application requiring flexible scheduling. In particular, the applicability of the framework to different application environments should be checked. These environments sould include industrial control systems, media processing applications, automotive embedded systems, and telecommunications. The output of the meeting will be a set of conclusions on the evaluation of the framework, which will set the work agenda on the subject for the following period.

Meeting Specific Research Meetings NXP - TUKL

Meeting Title	Specific Research Meeting NXP - TUKL
Date	TBD
Objectives and expected output	Multiresource management issues.

4.3 Cluster Integration: QoS Aware Components

4.3.1 Year 1 Achievements: Sept 2004 – August 2005

The main result of this work was the concrete identification of the integration topics and the start of this work. 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, and 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 will support (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 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 this period.

4.3.2 Year 2 Achievements: Sept 2005 – August 2006

The work has been aligned in three main activities:

- Specification of QoS properties using UML profiles and aspect-based approaches
- Generation of analysable models from the UML models
- QoS support in run-time components frameworks

INRIA has designed a technique to express the semantics of quality of service in a way compatible with classical functional and behavioural properties. The language for specifying quality of service relies on well-accepted and well-defined structures of the UML notation. Quality of Service properties (extra-functional issues, e.g. time, throughput, memory usage) are specified separately from the functional ones (classical types and pre/postconditions), using statecharts.

The transitions of these statecharts carry annotations that describe conditions of evolution as well as the side effects of this evolution on quality of service parameters. The composition of extra-functional statecharts with more classical, functional ones provides a specification that includes all aspects of the specification. QoS properties are thus handled using the separation of concern approach. QoS statecharts are in turn attached to software components, thereby extending the component specification with extra-functional properties.

The composition of extra-functional properties is managed by mapping extra-functional properties to functional ones on the components that implement these extra-functional properties. For instance, memory consumption of given components A and B is mapped to a memory abstraction component, which defines how A's and B's needs are satisfied and how they interact (e.g. whether the total amount required is below a threshold). The Inria technique leverages the concept of relativity of functionality: properties may be extra-functional under some points of view and purely functional under other points of view.

UPM has continued its work on the specification of QoS properties using the UML profile on "QoS and Fault tolerance". It has developed a set of guidelines to make it easier its use. UPM has also worked on the UML modelling of safety and time properties. In particular, UPM has worked on the modelling of safety properties is based on the safety evaluation processes from THALES ATM and EUROCONTROL. Then, it was defined a conceptual model that allows the representation of the safety properties, rules and methods. The next step was to develop annotations for UML that allows the description of the mentioned safety concepts. This work is being done and two approaches are being explored: use of the QoS and FT profile and using a newly created safety profile.

The main difficulties related with this work are the clear interpretation of the safety analysis process, the identification of the suitable elements in the conceptual model and the selection of the suitable UM L mechanisms for the modelling.

UPM has also experimented the generation of safety analysis models out of the UML annotated models. In particular, there are a couple of prototypes for fault-tree analysis and FMECA. The generated models can be used as input for commercial analysis tools (http://www.dit.upm.es/str , http://www.modelware-ist.org).

As a result of the ARTIST2 Workshop: "Beyond AUTOSAR" an effort was made at CEA, in collaboration with the Univesidad de Cantabria to address the identified necessity to perform evaluation of QoS and timing requirements of the component based applications. For this it is necessary to obtain as much as possible information of the components implementation from the manufacturers but preserving the intellectual properties involved. The idea is to ask the hardware/software components manufacturers to somehow partially "grey" their black boxes, requiring them to bring a timing behavioural model of the components that will be delivered, as soon as possible in the integration process. These models should be an abstraction of the functional behaviour implemented in the binary code, which will made explicit all the control flow paths for the worst case situations, the internal interactions, the mutual exclusive shared resources (critical sections), and the execution time consumed by each independent segment of code. Specific configuration and deployment services are required to be able to characterize and complete the models on the concrete platform to use. Once defined and integrated, these models will serve to evaluate response times as well as performance characteristics from a system wide viewpoint. The proposed strategy was presented in the 9th International Conference on Software Reuse 2006.

UPM and UC3M have worked on the improvement of the QoS management in the Robocop framework (http://www.hitech-projects.com/euprojects/space4u/index.htm, http://www.dit.upm.es/str). The overall design has been modified in order to make it easier the handling of QoS properties. UPM has collaborated in the standardization process of the MPEG Multimedia Middleware (M3W). The QoS management APIs are part of the standard (http://www.chiariglione.org/mpeg/technologies/mpe-m3w/index.htm).

CEA and THALES are working on container extension mechanisms within the COMPARE project in order to add the QoS specification for container services by providing examples from our prototypes. They have also participated on the development of the OMG QoS4CCM specification that can be downloaded from the OMG page (http://www.omg.org/cgi-bin/doc?ptc/2006-04-15).

4.3.3 Objectives and Work Planned: Sept 2006 – February 2008

The work on this activity will continue the most important topics identified, that are the development 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 analysis models for the QoS properties modelled.
- Composition mechanisms for determining whether it is feasible components interconnection and for deriving the non-functional characteristics of a group of connected components and the resources needed to fulfil them.
- Component frameworks to support the runtime composition of QoS aware components and to interact with the QoS management subsystems.

One open problem is how to include the description of extra-functional or QoS properties in functional models. There are results on this are such as two UML profiles with this aim: MARTE and "QoS and FT". There is some overlapping in the types of properties that can be described with them. In addition, there is little experience on their use. The goal is to experiment with them, using a suitable case study and QoS properties. The final results will be some guidelines for their use and an in-depth evaluation of both approaches. The use of other techniques, such a proposal based on aspects specification will be also explored. The integration of these frameworks is another topic of interest. CEA and INRIA have started an activity with the aim of integrating their research results in this area.

One of the advantages of the modelling of QoS properties is to check a system design meets the QoS requirements. For this purpose, the desirable approach is to be able to extract the information with respect to a QoS property and generate a model that can be used for analysis purposes. There are some problems related with this aim, such as to make sure that the QoS information provided is the appropriate and sufficient for the analysis and to evaluate the suitability of the different modelling approaches for this purpose.

On relation with the composition of components, the goal is to provide methods to assemble relevant quality properties and resource requirements. There are actions to develop contract models where the provided and required qualities between components are considered and ensured. These operations can be done at design time or at run time. In this case, it is needed to consider platform dependent transformations for including resource usage information to the quality levels that a component can provide.

It is clear that the problems are not only at the development phases. Another problem to handle is QoS component frameworks, which allows at runtime the identification, retrieving, consistency checking, registering and composition of components. The QoS properties of the components need to be taken into account to ensure a proper system operation. For example, when trying to identify a suitable component for an application, it is necessary to make sure that in addition of the required functionality, it is able to provide a certain QoS with respect to some meaningful properties. There are some works in the group related with this topic. The goal will be to try to evaluate the relative merits of the runtime frameworks developed by the partners (UPM, THALES and CEA) and to propose an API that integrates the functionality required for this work and to try to promote it. CEA and THALES will use the connector extension of CCM to add fault tolerance properties.

4.3.4 Meetings Planned

Meetings for Year 3 will be decided at the annual consortium meeting, Nov 7th in Paris (just before the review).

4.4 Real-Time Languages

4.4.1 Year 1 Achievements: Sept 2004 – August 2005

This activity was not active during this period.

4.4.2 Year 2 Achievements: Sept 2005 – August 2006

This activity started in the middle of Year 2.

Work on Ada 2005

A number of sites have started to evaluate the new features available in Ada 2005. A useful meeting was held in March (see description below).

The university of York has begun to build a library of reusable facilities to ease the programming of flexible real-time systems. This has concentrated on the programming of standard patterns such as periodic tasks (with deadline and CPU overrun detection), and sporadic tasks (with deadline, CPU time and minimal separation violation detection).

The University of Cantabria has been working on the integration of the real-time services built in the MaRTE OS kernel with the GNAT Ada compiler system, with the goal of generating a platform that fully supports the new real-time features defined in the Ada 2005 standard. This platform will be used by the other partners involved in this activity to integrate further research in the area of real-time languages, and will be made available to industry and the Ada community in general as free software.

The particular activities carried out in this period at the University of Cantabria have been the addition of the timed event service to the MaRTE OS kernel, the integration of this facility with the GNAT run-time system, and the study of the changes required in the GNAT compilation system for supporting Ada 2005 real-time services. In addition, a development environment has been setup for the integration of the compiler run-time system with the kernel, and the entire Ada official test suite has been run in order to validate the new platform.

Work on other languages

Ada has in many ways lead the way in attempting to bring into language design the abstractions and idea that have been developed within the real-time research community. Other languages have are also moving in that direction. The C language has been closely linked to POSIX and ARTIST is participating in POSIX standardisation work (although this is not a very active area at the moment). Java, through its RTSJ (real-time specification for Java), has many of the flexibility features desirable in an implementation language – again ARTIST is fully involved in the definition work surround RTSJ. ARTIST is therefore sponsoring and co-organising JTRES (see future work). The work that has been undertaken on Java and Ada has also been used to influence the develop of concurrent and real-time versions of Eiffel (SCOOP). Again ARTIST sponsored a workshop on this language topic.

Other work has been focused on research languages. AbsInt participates in the ST-project Embounded (www.embounded.org). Within Embounded a new functional programming language (Hume) for real-time applications is under consideration. The aims of the EmBounded Project are to identify, to quantify and to certify resource-bounded code in Hume, a domain-specific high-level programming language for real-time embedded systems. Using formal models of resource consumption as a basis, the project will develop static analyses for time and space consumption and assess these against realistic applications for embedded systems.

In almost all real-time systems, the chosen implementation language, the software synthesis method and/or the coding guidelines have a strong influence on the analyzability of real-time systems. Disciplined code as synthesized from specification languages like SCADE, ASCET-SD and others allow the determination of safe and highly precise execution-time bounds. Undisciplined code and the use of many dynamic languages features, e.g. dynamic method dispatch or dynamically allocated heap data structures, lead to high degrees of overestimation. It follows that there is a link between the work of this activity and that undertaken on WCET elsewhere in ARTIST2. This link will continue to be explored.

4.4.3 Objectives and Work Planned: Sept 2006 – February 2008

A number of workshops are planned for the coming period, including the next International Real-Time Ada Workshop (IRTAW13) which will take place in the US (Boston area) in March 2007. This will be sponsored by ARTIST and involve a number of ARTIST participants (see earlier discussion). ARTIST is also co-organising JTRES – 4th Workshop on Java Technology for Real-time and Embedded Systems. This is scheduled for 11-13 Oct 2006 in Paris. See: http://www.artist-embedded.org/artist/-JTRES-2006-Java-Technologies-for-.html

Work will continue on the evaluation of Ada 2005 including the use of new Ada 2005 real-time capabilities for EDF-based server scheduling (including capacity sharing and capacity stealing servers), and the use of proof-carrying code in Embedded Systems. Much of this work will be reported to IRTAW13. A further ARTIST meeting to look at implementation issue will occur in Spain in October 2006.

Both of the above events will be used to define and extend the pattern of use that will allow these languages to be used in the programming of effective, adaptive and flexible real-time systems.

This activity will also continue to act as a coordinator for a wide range of language work going on within Europe.

4.4.4 Meetings Planned

Meetings for Year 3 will be decided at the annual consortium meeting, Nov 7th in Paris (just before the review).

5. Cluster: Compilers and Timing Analysis

Cluster Leader: Reinhard Wilhelm (Saarland)

The following is a description of the activities and overall objectives for the period: September 2006 – February 2008. The next reporting period will cover September 2006 – August 2007.

Staff Mobility

Mälardalen had two MSc students from Univ. des Saarlandes visiing for one week in January 2006. Also, Raimund Kirner from TU Wien visited us for a week in the fall 2005 (Oct 30 - Nov 5).

5.1 Platform: Timing Analysis Platform

5.1.1 Year 1 Achievements: Sept 2004 – August 2005

CRL2 has been chosen as intermediate representation language to be used among partners. An efficient C/C++ library for CRL2 is created, providing API to data structures. Parts of analysis tool chain communicate via CRL2. AbsInt started to develop definition and documentation of external text format for CRL2.

Several industrial case studies have been performed by students at Mälardalen.

Integration of Bound-T with Mälardalen flow analysis.

5.1.2 Year 2 Achievements: Sept 2005 – August 2006

A detailed description of these achievements is provided in the activity's deliverable.

- **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 Interchange') 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 Mälardalen approach.
  In the joint approach we perform a parametric WCET analysis based on the aiT-toolchain. 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/docdetail.php?DID=1975&viewmode=paper&year=2006

5.1.3 Objectives and Work Planned: Sept 2006 – February 2008

Most of the work will concern the basis for tool integration, the chosen interface language AIR and its extensions for different purposes.

1. **AIR Semantics**
   The semantics of the chosen interface language will be specified.
2. **Computation semantics language**
   a. Discussion of the computation semantics language as an extension of AIR.
   b. Definition of the primitive operations
   c. Specification of the semantics for a core language
3. **ALF flow analysis input format**
   The ALF flow analysis input format will be finalized, and Mälardalen's flow analysis will be adapted to use ALF instead of the current NIC format.
4. **The WCET Tool Challenge**
   The WCET Tool Challenge will be run in the fall 2006. It is based on a set of WCET benchmarks. Developers are invited to submit their WCET tools to be run against these benchmarks. The purpose is 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 benchmarks. (cf. http://www.idt.mdh.se/personal/jgn/challenge/). This is only the first instance of the WCET Challenge, which is intended to be an annual event.
5. **Measurement-Based Analysis**
   The work on coverage criteria for the measurement-based analysis will be continued and documented in a report. Further, the prototype framework for test-data generation will be improved.

5.1.4 Meetings Planned

Meetings for Year 3 will be decided at the annual consortium meeting, Nov 7th in Paris (just before the review).

5.2 Platform: Compilers Platform

5.2.1 Year 1 Achievements: Sept 2004 – August 2005

5.2.1.1 From the previous “Compilers Platform” Activity

Members of the compilers platform activity have evaluated different options for a common compiler platform that can be shared for most intra-cluster activities. After a review of these options, including platforms like gcc, lcc, SUIF, and OCE, it was concluded that the CoSy platform from ACE will be selected as the primary platform, due to its flexibility and robustness. However, due to the partners´ specific preferences and interests, also other secondary platforms (ROSE, ICD-C) will be used.

A dedicated free CoSy research license for ARTIST2 partners has been agreed with ACE. ACE has supported the cluster teams, especially at ST and Aachen, in their use of the CoSy compiler development system by providing advice on its use as well as modification and extensions. Periodic face-to-face meetings, including two global cluster meetings, helped facilitate this.

Furthermore, the compilers cluster has been organized into a set of mini-clusters with 2-3 parties each that focus on specific research and integration aspects.

A number of discussions were held between AbsInt and ACE to consider the design issues involved in CoSy-PAG integration. The project has been put on hold until sufficient commercial interest reasserts itself as, despite the inherent interest of ACE and AbsInt in the work, engineering more than a point solution or proof of concept requires more resource than the companies are able to justify at present.

TU Vienna worked on development of the PAG Interface Generator (PIG) to simplify the integration of PAG into existing infrastructures. PIG was successfully used to integrate PAG into different C/C++ infrastructures. The integration in ROSE covers all C language features. For testing the PAG integration a constant propagation for C was specified and evaluated for different infrastructures.

5.2.1.2 From the previous “Architecture-Aware Compilation” Activity

Cluster members have been studying compiler requirements for modern embedded processor architectures. As a result, several specific code optimizations have been identified that need to be supported, e.g. SIMD instructions, fast local memories, as well as reconfigurable architectures.

The compilers cluster has then been organized into a set of mini-clusters with 2-3 parties each that focus on code optimization aspects. Alpha versions and software prototypes have become available (see year 1 deliverables).

The common compiler platform CoSy (see 3.2) has been supported whenever possible. Furthermore, tight links to the Execution Platforms cluster, specifically with Bologna, have been established.

5.2.2 Year 2 Achievements: Sept 2005 – August 2006

5.2.2.1 From the previous "Compilers Platform" Activity

A detailed description of these achievements is provided in the activity's deliverable.

Design of a WCET-aware C Compiler (Dortmund – Absint)
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 (Dortmund – Absint)
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

Optimisation of Conditional Execution in CoSy (ACE – Aachen)
A dynamic programming algorithm is being implemented and tested on a number of different architectures to validate its behaviour with real world code and current high-end industrial processors.

A prototype comprising a set of optimisation engines and compiler has been constructed.

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

Extension of the ROSE-PAG integration from C to C++ and Implementation of Alias Analysis. (Absint – TU Vienna)
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
(Absint – TU Vienna)
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).

Agreement on Basic Guidelines (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.

Work on Automated Checkers (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. As a starting point, the verification and development of checkers for loop transformations based on unimodular transformations is investigated.

5.2.2.2 From the previous “Architecture-Aware Compilation” Activity

A detailed description of these achievements is provided in the activity’s deliverable.

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.

Dortmund-Aachen cooperation on Bit-True Data Flow Optimizations as a Processor specific Source-Level Transformation: During the last reporting period, it was found bit-true dataflow analysis and the corresponding SIMD optimizations implemented at the C source-level serve greater benefit than as a post-pass analysis into the LISATek tools, as they can be easily retargeted for different processor architecture. Currently, the data flow analysis along with three different optimizations [Fal06] is implemented for TI C6x DSP. The first optimization detects and optimizes saturated arithmetic operations. The second optimization looks for SIMD instructions for packed parallel arithmetic. Third, is strength reduction optimization which determines the number of unused bits in variables and then reduces them to the smallest data type.

ACE - Aachen Cooperation:

SIMD Support in CoSy
Aachen and ACE are working on an extensible framework to allow compiler writers to target SIMD hardware more quickly and efficiently. At present, hand optimised point solutions tend to be used in industry which are not conducive to retargeting. Such an approach is essential as part of an overall solution to support SIMD hardware generated from architectural descriptions
One-line Description of the Outcome
Work in progress – extensions to data dependency analysis, support for interprocedural pointer alignment analysis and code generation description extensions to enable SIMD retargetability.

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

One-line Description of the Outcome
A prototype comprising a set of optimisation engines and compiler has been constructed.

AbsInt - TU Vienna Cooperation:

Integration of PAG in the ROSE C++ Infrastructure and Evaluation of C++ programming styles

The ROSE-PAG integration achieved in Y1 for C was substantially extended to cover full C++ (only exlcuding Exceptions). This includes handling of templates, virtual methods, short-circuit evaluation in conditions, resolving overloaded functions, C++ namespaces, constructor and destructor calls. Based on that infrastructure an initial version of a tool for whole-program analysis (WHOPA) was implemented for performing high-level evaluation of different generic programming styles suitable for embedded systems.

Evaluation of C++ optimizations
The high-level analysis of the C++ evaluation cases shows a significant difference in code size after template instantiation and in-lining (which is crucial for the relevant codes to permit whole program optimization). The evaluation cases, although performing the same operations on test data, show a difference by a factor of six in over all codes size including library codes.

The generated assembly code was measured for different optimization levels and additionally evaluated by hand. Our early results show that a certain class of generic programming styles can be automatically optimized such that we obtain similar assembly codes as with low-level C or assembler programming; this indicates that not only C but even very high-level C++ techniques might become suitable for programming embedded systems in near future.

STMicroelectronics: Compiler for reconfigurable processors
NB: STMicroelectronics is leaving the NoE at the end of Year 2.
During year two, more work was achieved on the retargetability of the compiler. It has been mainly devoted to the continuation of the effort to make the process fully dynamic. It includes:

- the completion of the definition of the scope of the possible extensions,
- the definition and implementation of the modification to be made in the compiler toolchain to make its reconfiguration fully dynamic,

- the completion of the work on the model to be built to allow the configuration of each toolchain component. The scope of the reconfiguration toolkit that was already in place to configure assembler and linker, was thus enlarged,
- the definition of the interface and tools needed to allow end users with different levels of expertise to defined their own extensions. The implementation of the end-user graphical interface has been outsourced to an external company (Coware/LISATek).

Longer term perspectives have been anticipated: for instance, we already know that SIMD instructions is a very demanded features in most recent applications and SOC. Some of the choices were made to ease automatic vectorization of SIMD extensions code in the future.

5.2.3 Objectives and Work Planned: Sept 2006 – February 2008

5.2.3.1 From the previous “Compilers Platform” Activity

The activity will focus on further integration and refinement the compiler platform(s). It is foreseen to retain the existing “mini-cluster” structure outlined in the Compiler and Timing Analysis Cluster report and its collaboration scheme.

University of Dortmund
The development of WCET-aware compiler optimizations will be performed by Dortmund University. On the one hand, this will include optimizations exclusively focussing on WCET as objective function, like e.g. exploitation of memory hierarchies for WCET minimization. On the other hand, the mechanisms provided by Dortmund’s WCET-aware compiler developed during ARTIST Year2 for multi-objective optimization (e.g. trading off WCET vs. code size) will be used. It is intended to set up a cooperation with ARTIST2 core partners of the timing analysis platform (Mälardalen, Vienna) in order to integrate flow analysis techniques into Dortmund’s compiler.

IMEC
For the next 18 months during the collaboration between Dortmund Uni and IMEC vzw, it is planned to develop source-to-source optimization techniques mainly for statically allocated data, dealing with arrays inside loop bodies, but will also support dynamically allocated data like linked lists. The techniques will exploit certain properties of scratchpad or cache based memory hierarchies and will ensure data coherency in the memory subsystem through integration of timinig analysis and compilers.

Absint
Recent developments in the field of register allocation allow for seperation of its subproblems thereby enabling better solution techniques. In this context Saarland University investigates coalescing, a compiler back-end optimization improving performance and reducing code size.

The solution technique relies on the strict separation of the subtasks and uses the general purpose optimization technique Integer Linear Programming (ILP) to derive (near-)optimal results. These are important, because in embedded systems the size of software is direclty related to production costs. To make this ILP-approach feasible, several optimizations were implemented improving the performance of the solving process to reasonable times. First experiments look promising and finial results in terms of speed up or code compaction of the compiled code will be available soon. This work is partially supported by the German Research Foundation (DFG) through the Graduiertenkolleg 623.

Technical University of Berlin (new core partner)
Berlin's work, in cooperation with ACE, will focus on transformation of loops and their verification. It is planned to develop checkers for general unimodular loop transformations to ensure that, in individual compilations, the source and target programs are semantically equivalent.

ACE – Aachen
ACE will continue to support partners in using the CoSy compiler platform. Aachen will support and partially coordinate these activities and will provide partners with further support on the CoSy Express and LISATek technology. Furthermore, ACE and Aachen will continue their tight cooperation on platform based code optimization engines.

AbsInt – TU Vienna Cooperation
TU Vienna will further enhance the ROSE-PAG connection for performing whole-program source-code analysis of C/C++ applications and provide evaluation data on the scalability of an analysis. The goal is to provide enabling technology for library centric development of embedded systems codes. AbsInt and TU Vienna will continue their cooperation on automating the integration of PAG in existing platforms and in developing program analyses with PAG.

5.2.3.2 From the previous "Architecture-Aware Compilation" Activity

The activity will focus on further integration and refinement of architecture aware compilation technologies based on the compiler platform(s). It is foreseen to retain the existing "mini-cluster" structure outlined in the Compiler and Timing Analysis Cluster report and its collaboration scheme.

For the next 18 months during the collaboration between Dortmund Uni and IMEC vzw, it is planned to develop source-to-source optimization techniques mainly for statically allocated data, dealing with arrays inside loop bodies, but will also support dynamically allocated data like linked lists. The techniques will exploit certain properties of scratchpad or cache based memory hierarchies and will ensure data coherency in the memory subsystem through integration of timinig analysis and compilers.

ACE and Aachen's work on SIMD will continue to improve its ability to translate a wider class of loops into SIMD instructions. The SIMD and conditional execution work has to be integrated into CoSy solving various practical retargeting issues.

AbsInt – TU Vienna Cooperation: Perform whole-program source-code analysis of C/C++ applications and provide techniques and evaluation data for scalable analysis. The goal is to enable library centric development of embedded systems codes.

5.2.4 Meetings Planned

Meetings for Year 3 will be decided at the annual consortium meeting, Nov 7th in Paris (just before the review).

5.3 Cluster Integration: Architecture-aware compilation

5.3.1 Year 1 Achievements: Sept 2004 – August 2005

5.3.2 Year 2 Achievements: Sept 2005 – August 2006

5.3.3 Objectives and Work Planned: Sept 2006 – February 2008

5.3.4 Meetings Planned

Meetings for Year 3 will be decided at the annual consortium meeting, Nov 7th in Paris (just before the review).

6. Cluster: Execution Platforms

Cluster Leader: Lothar Thiele (ETHZ)

The following is a description of the activities and overall objectives for the period: September 2006 – February 2008. The next reporting period will cover September 2006 – August 2007.

Staff Mobility

Area of Collaboration:	**Low power real-time systems for multimedia applications**.
Clusters:	Execution Platforms, Compilers and Timing Analysis
Sending Partner:	Peter Marwedel – Dortmund Receiving Partner: Petru Eles - ESLAB/Linköping
Person:	Olivera Jovanovic
Technical Work:	Extend current techniques for on-line voltage/frequency scaling to multiprocessor systems.
Dates	April - November 2006
Published Work	Journal publication planned

6.1 Platform: System Modelling Infrastructure

6.1.1 Year 1 Achievements: Sept 2004 – August 2005

The work during the first 12 months was focused on two different approaches to system modelling:

- Simulation-based modelling, with the aim to integrate cycle-true models with transaction level and abstract models in order to be able to do cross-layer and –level modelling and analysis.
- Formal modelling, with the aim to integrate different formalisms in order to be able to perform worst-case and response time analysis as well as schedulability analysis.

The main results from the simulation-based modelling are:

- University of Bologna established a consistent SystemC based and cycle accurate simulation environment for on-chip multi-processor systems, with software support and hardware extensions for multi-processing (synchronization, inter-core communication, optimized transfer engines, etc.) targeting the embedded computing domain.
- Technical University of Denmark established a system-level simulation framework based on SystemC which allows to model and simulate cross-layer dependencies between application software, RTOS middleware and platform architecture, including processors, memories and interconnect structures.

- University of Bologna and Technical University of Denmark developed and demonstrated an advanced traffic generator model, featuring extensive reactive capabilities to mimic the behaviour of the core when facing unpredictable environmental events and network performance, and awareness of the multiprocessor nature of the target platforms, which implies synchronization requirements.

The main results from the formal modelling are:

- Technical University of Braunschweig made considerable extensions to the SymTA/S tool to model and analyse power consumption of complex heterogeneous embedded systems. The task model was extended to take multiple remote transactions during the task's execution into account during response time analysis. Finally, sensitivity analysis was introduced into the model. Sensitivity analysis allows the designer to determine the maximum head room achievable for each component in the system with respect to various system performance properties.

- University of Linköping developed modelling and analysis methods for heterogeneous distributed embedded systems within automotive applications. The methods are based on formal techniques, allowing for a holistic schedulability analysis.

6.1.2 Year 2 Achievements: Sept 2005 – August 2006

A detailed description of these achievements is provided in the activity's deliverable.

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. Affiliated 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.

6.1.3 Objectives and Work Planned: Sept 2006 – February 2008

In the next 18 months, we will continue our effort in developing and refining the various system models. We will continue our model integration both within and between the simulation-based and formal-based modelling approachs. The model refinment will be based on the feedback gained by using the models, in particular with the activities of Communication Centric Systems and Low Power Design.

TU Braunschweig will continue its work extending the semantic model of SymTA/S to efficiently cover MPSoC architectures. Additionally, TU Braunschweig will conduct further research in sensitivity analysis techniques and its application to predictable system design.

University of Linköping will continue it development on the simulation environment based on the ARTS framework from Technical University of Denmark. They will conduct experimental evaluations using the simulation environment for distributed embedded systems. The focus is on evaluation of the impact of various protocols on worst case and average performance; evaluation of pessimism of various response time analysis approaches; impact on quality of control.

University of Bologna will investigate techniques and approaches to reduce simulation time. This is essential for analysis of complex platforms and of complex workload, in particular when dealing with cycle-accurate models. To this purposes simulation-acceleration techniques base on hardware emulation will be investigated.

Technical University of Denmark will continue research on the ARTS environment. Extensions will cover; modelling capabilities for dynamically reconfigurable architectures. This requires that not only the software can be moved and modified during platform execution, but also the hardware itself. The aim is to be able to model and analyse new architectures for reconfigurable computing. The research on modelling wireless sensor networks will be continued. Furthermore, extension of the modelling capabilities toward lab-on-a-chip will be started, in particular towards biochips, i.e. platforms which are able to move microfluidic droplets around within an array of cells in order to mix and analyse chemical liquids.

Technical University of Denmark will continue the work on linking simulation models with formal models. In particular they will extend their effort in formalizing the ARTS model using timed automaton based on UPPAAL. This will allow the same platform model to be expresed as a simulation model and as a formal model. The work will be carried out in cooperation with the research group at CISS in Aalborg which is a partner in the cluster on test and verification.

ETH Zurich will combine simulation and analytic methods: Continuing the work with University of Bologna and possibility strengthening the cooperating with Technical University of Denmark on combining simulation and analytic estimation methods.
The main approach is to use simulation and dedicated benchmark applications in order to profile a hardware platform with respect to OS and communication overhead. Using this information to parameterize analytic performance analysis approaches.

6.1.4 Meetings Planned

A meeting of the Execution platforms cluster will be held May 14th in Linköping.

6.2 Cluster Integration: Communication-centric systems

6.2.1 Year 1 Achievements: Sept 2004 – August 2005

The most important results during the first 18 month were achieved through the integration work of the activity partners.

First, the stat-of-the-art in modelling and performance verification was assessed:

- ETH Zurich and the Embedded Systems Institute Eindhoven investigated how different performance verification approached can be integrated into the system design process. In a case study different approaches were utilized for analysis in order to identify their strengths and limitations in the design process.
- Technical University of Braunschweig and University of Linköping compared and evaluated so-called holistic and compositional performance analysis approaches. Also, methodologies for the analysis of complex hierarchical and dynamic priority schedulers were investigated.

The results of the comparisons and case studies were considered for the creation of new powerful techniques and models for the performance evaluation and optimization of complex real-time systems:

- ETH Zurich and University of Bologna utilized formal performance analysis to speed up cycle accurate simulation.
- University of Bologna and University of Linköping established high-level models for shared communication.
- Technical University of Denmark and University of Bologna created a methodology for mixed level simulation allowing an efficient evaluation and optimization of NoC architectures.
- Technical University of Braunschweig and ETH Zurich investigated a mixed performance analysis approach using SymTA/S and Real-Time Calculus.
- University of Braunschweig and University of Notre Dame developed a power analysis for complex heterogeneous embedded systems, which is currently used for power optimization.

6.2.2 Year 2 Achievements: Sept 2005 – August 2006

A detailed description of these achievements is provided in the activity's deliverable.

Timing Analysis of the Flexray Protocol (University Linköping)
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 (University Linköping)
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 (ETH Zürich)
Collaboration with Arne Hamann, University of Braunschweig
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 basic 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 (ETH Zürich)
Collaboration with Marcel Verhoef at Chess Information Technology, Embedded Systems Institute Eindhoven and Radboud University Nijmegen, and with Paul Lieverse at Siemens VDO
In this activity, we 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 (TU Braunschweig)
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, analysing and modelling real-time aspects. To be able to use different tools specialized in handling real-time 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 (TU Braunschweig)
Cooperation with Sharon Hu and Bren Mochocki from University of Notre Dame
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 (TU Braunschweig
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 (TU Braunschweig)
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 DDR-SDRAM frame memory bandwidth. This design was awarded with the "DATE2006 Design Record" distinction [LHR+06].

Simulation-based analysis of SoC interconnection architectures (University of Bologna)
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 (Technical University of Denmark)
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 becomming increasingly difficult. MANGO employs virtual channels to provide connection-less best-effort routing as well as connection-oriented virtual circuits, for which service guarantees can be given. The predictability of guaranteed services is a way to promote system-level integrity. The MANGO architecture has been demonstrated through a circuit-level design of a 5x5 router using a 0.13 μm CMOS standard cell library from STMicroelectronics. Netlist simulations showed a performance of 650 Mflits/s under typical timing conditions [BS06]. Three patents [BS05] on the MANGO technology have been filed and a startup company, called Teklatech (www.teklatech.com), was formed as a spin-off from this research. Teklatech is developing a one-step EDA solution to achieving timing closure in large scale, globally synchronous, deep submicron ASIC designs.

Distributed wireless sensor networks (Technical University of Denmark)
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 build. 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 needs 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 (Technical University of Denmark and University of Bologna)
In Multi-Processor System-on-Chip (MPSoC) design stages, accurate modeling of IP behaviour 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.

6.2.3 Objectives and Work Planned: Sept 2006 – February 2008

TU Braunschweig and ETH Zürich

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.

ESLAB Linköping

1. Further development of optimization approaches for systems built on heterogeneous communication protocols, in particular, Flexray.
2. Further development of the analysis and optimization techniques for fault-tolerant distributed systems.

These techniques will be incorporated into the tools developed by the various partners, in particular, Symta/S in Braunschweig.

TU Braunschweig

1. Further research in the extension of the SymTA/S model with hierarchical event models.
2. Development of advanced techniques for system robustness optimization. This would leverage the applicability of the SymTA/S methodology for reliable and predictable embedded system design.
3. Development of mapping optimization approaches with automatic communication synthesis.
4. Refinement and further development of semantical extensions for formal analysis of MPSoCs with focus on shared memory accesses. This includes the coupling of the tools SymTA/S und SymTA/P (both developed at TU Braunschweig). This activity could profit from synergy effects with 1.

University of Bologna

1. Network-on-chip architecture exploration: a more modular and scalable system-inteconnect architecture development approach will be studied in details, along with cross-benchmarking against traditional system interconnects

Technical University of Denmark

1. Further development of network-on-chip architectures and exploration at both circuit and system level. Development of NoC benchmarks and conducting comparative NoC studies.
2. Further development of communication models for distributed embedded systems, in particular wireless sensor networks and fault-tolerant distributed systems.
3. Development of mapping approaches which explores optimized communication architectures in terms of metrics like performance, power consumption, cost and fault-tolerance.

6.2.4 Meetings Planned

A meeting of the Execution platforms cluster will be held May 14th in Linköping.

6.3 Cluster Integration: Low-Power design

6.3.1 Year 1 Achievements: Sept 2004 – August 2005

The main contribution of University of Bologna has been in the extension of a complete power-modelling infrastructure for all components of current multi-processor systems-on-chip platforms and for future Network-on-Chip-based platform. Several extensions 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 affiliated partner STMicroelectronics). Techniques for energy optimization in system interconnects have been explored with the help of this platform

Additionally, Bologna has started a research effort on energy aware mapping of multi-task applications on multi-processor SoC execution platforms. 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.

The Technical University of Denmark has started the development of a sensor network platform (Hogthrob project), with focus on: (1) Low Power processor design based on low-power synthesis (e.g., clock-gating), power modes and de-synchronizing (2) Power modelling: Simulation-based power modelling and estimation techniques have been investigated, with emphasis on stochastic modelling of batteries and investigation into the macro-modelling of various hardware components.

Additionally, DTU has also worked on empirical power estimation: 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.

Linköping University has developed a technique for static routing on NoC, with guaranteed delays and arrival probabilities in the presence of transient faults. 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 has also performed additional work aiming at a more accurate modelling of actual communication and memory techniques used in MP SoC. 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.

6.3.2 Year 2 Achievements: Sept 2005 – August 2006

A detailed description of these achievements is provided in the activity's deliverable.

Power modeling 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 affiliated 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 modeled and simulated at a very high level of accuracy (cycle-accuracy and bus-signal-accuracy), and compared with mixed AHB/AXI platforms

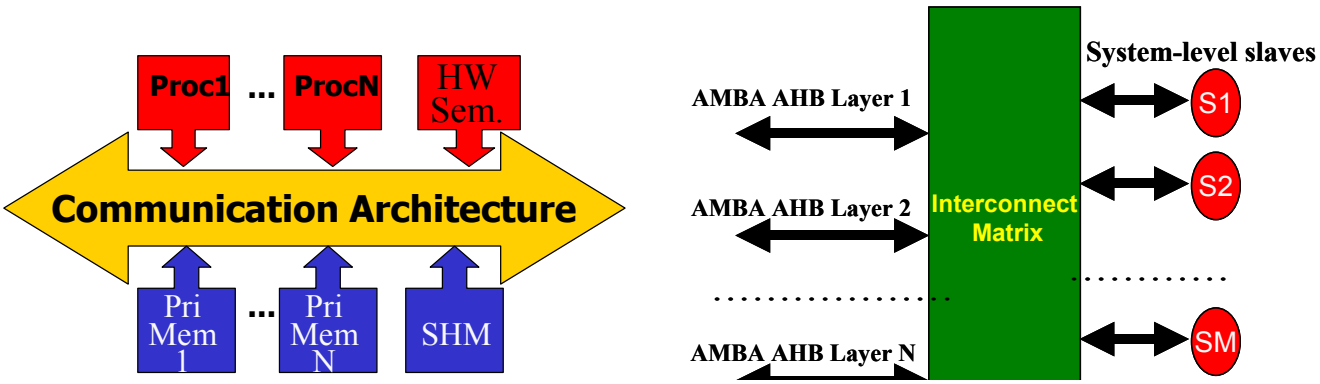

Fig 1 (a) Baseline single-node shared bus platform (b) advanced multi-layer interconnects

The original MPARM platform allowed the modeling and simulation of single-node communication architectures (as depicted in Fig.1a). The platform was enhanced with the possibility to extend the modeling 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 favoring 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 modeling 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 behavior 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.

Outcome

The outcomes of this activity are: the development of a virtual platform for power modeling of complex multi-core systems on chip. This platform will facilitate further integration among partners and associates, thanks to is flexibility and generality.

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 behavior.

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.

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. We have 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.

Outcome
The outcome of this activity is the development of a methodology for design-time allocation, scheduling, frequency and voltage setting for multi-task applications onto MPSoC platforms. This outcome is a starting point for follow-up integration activities aiming at the extension of the methodology to more dynamic problems, where run-time decisions will be required

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.

Outcome
The outcome of this work is a novel scheduling strategy (called lazy scheduling) that is well suited to energy-harvesting systems operating under real-time constraints. It is the first result of this kind in this quickly growing research area and received a lot of attention in the scientific community. Two joint publications have been written.

6.3.3 Objectives and Work Planned: Sept 2006 – February 2008

Power optimization via system-level resource allocation and scheduling

Linkoeping and Bologna will continue cooperation on this topic, and we will explore the possibility of including the evolutionary exploration from DTU. Our main goal for the coming period is to consolidate the results regarding the optimization of energy-efficient time constrained multiprocessor systems. The main directions are the following:

- improve and refine the task-based application models as well as the architecture models, in order to make them as realistic as possible, in the context of current execution platforms and target applications;
- explore more efficient design space exploration approaches based on mathematical programming or heuristics, e.g. evolutionary algorithms;
- extend the approaches to on-line voltage scaling, such that dynamic slack can be exploited; the problem is particularly interesting in the context of multiprocessors;
- Explore interaction and tradeoffs between energy efficiency and fault tolerance.

Scheduling based energy optimization for energy-scavenging wireless sensor networks ETHZ and Bologna will continue to work on low power sensor network design. In particular, the results so far will be extended towards application-level decisions.

- To this end, on-line control strategies need to be developed that change the state of the application depending on the current systems state, e.g. the amount of local data stored, and on the estimation of the future energy flow. The approach will be based on the experience in UoB on building energy-harvesting nodes and on convex optimization from ETHZ. Both theoretical and practical aspects of the problem will be investigated. The feasibility of implementation of advanced on-line control strategies on tightly constrained sensor network hardware platforms will be explored. This activity will leverage the experience on ETHZ and UoB on the hardware-software design of wireless sensor nodes.

UoB and ETHZ are actively developing a joint prototype sensor node which is powered by solar energy. To this end, a BTnode - originally developed at ETH Zurich - has been transferred to Bologna. Equipped with solar panels and supercapacitors for energy storage, measurements on this prototype will illuminate the practical relevance of our theoretical results.

6.3.4 Meetings Planned

A meeting of the Execution platforms cluster will be held May 14th in Linköping.

6.4 NoE Integration: Resource-aware Design

6.4.1 Year 1 Achievements: Sept 2004 – August 2005

During the first six months of the project, several cooperations have been set up. Cooperation has been established between Università di Bologna and Dortmund University. The objective is to integrate the memory-aware compiler developed in Dortmund with the multi-processor platform simulator developed in Bologna. The first results have been development of a new source-level transformation tool for performing memory optimization by Dortmund, and the development of compatible memory organization models by Bologna (including I and D caches as well as scratchpad memories).

A second cooperation has been established between Università di Bologna and Aachen University. 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 has been the 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 has been established. The objective of this cooperation is the exploitation of the platform simulator developed by Bologna, and 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.

6.4.2 Year 2 Achievements: Sept 2005 – August 2006

A detailed description of these achievements is provided in the activity's deliverable.

- **Energy efficient time constrained systems**
  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 the efforts have concentrated on an extension of this approach to multiprocessor systems. This work is currently performed as part of the ARTIST mobility action in cooperation with the Dortmund group and will be continued into the following period.

- **Predictability in Multiprocessor SoC architectures**
  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 (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

- **Integration of LISATek ISS models in SystemC and the MPARM virtual platform**
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.

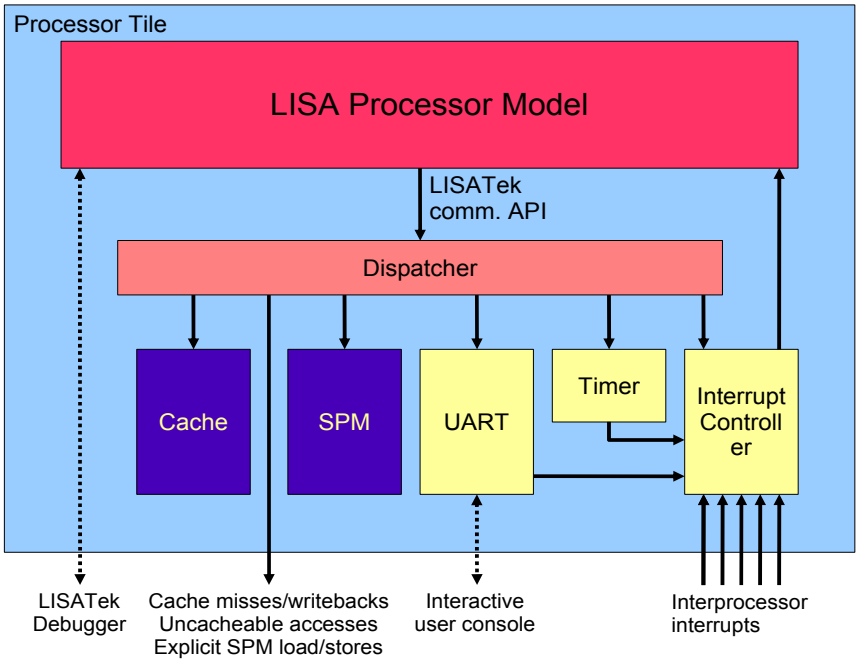

Figure 1 Processor Tile

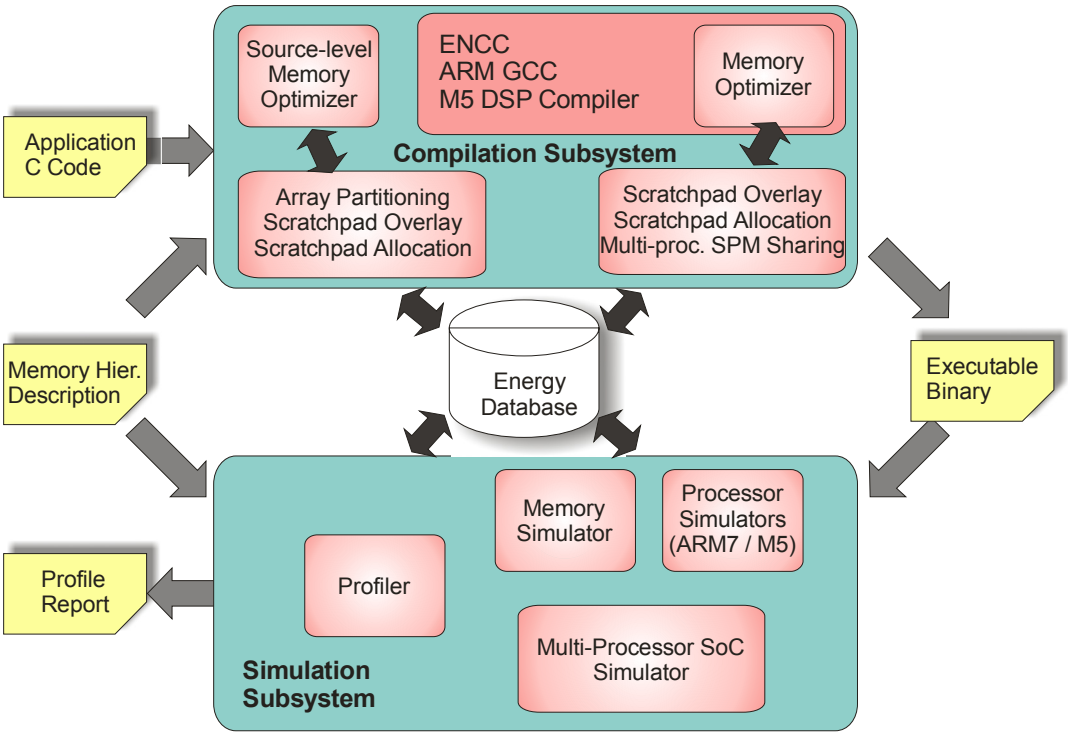

Figure 2: Memory Aware Compilation and Simulation Tool-Chain

- **Memory Aware Compilation and Simulation Tool-Chain for Energy Optimizations**
  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). 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**
  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**
  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

- **Analysis of cache predictability**
  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**
  The University of Saarbrücken and AbsInt (an industrial member of the compiler cluster) have cooperated with ETAS (an external company) 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].
  Web site: http://en.etasgroup.com/about/tradeshows/documents/2006-03-15_AutomotiveSoftwareWorkshop_ASCET_Paper_Renz.pdf

- **Resource aware design space exploration**
  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 has been accepted for publication at DIPES 2006 [Mad06]

- **FET Open Call project proposal**
  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. However, only 5% of the proposed projects can be funded, and this project will probably not be among them.

- **Interfaces for real-time components**
  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].
  Web site: http://chess.eecs.berkeley.edu/pubs/92.html
- **Resource awareness in sensor networks**
  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.

6.4.3 Objectives and Work Planned: Sept 2006 – February 2008

The Linköping group will continue to work in the context of predictability on multiprocessor SoC. An integrated environment will be implemented including WCET estimation, system scheduling and optimisation, taking into account the system wide interactions. Furthermore, the cooperation of the Linköping group with Braunschweig on WCET analysis and Bologna on the hardware implementation will continue. Cooperation with Dortmund on frequency and voltage scaling will continue.

Over the next reporting period, efficient utilization of memory hierarchies by multi-process applications will be the focus of the research at Dortmund. In the coming months, the group at Dortmund will enhance the memory allocation approaches to improve the scratchpad utilization by multi-process applications comprising of periodic and aperiodic tasks. A further improvement of the proposed approaches will be achieved by supporting applications with cooperative multi-tasking. The group also plans to develop memory allocation approaches as part of its compilation framework which is tightly integrated with the multi-processor SoC (MPARM) simulator from University of Bologna. The integration of the highly configurable memory hierarchy simulator developed at Dortmund into MPARM is also planned for the coming months. It is expected that these research efforts will further strengthen the cooperation between University of Dortmund and University of Bologna.

The group at the University of Saarbrücken will extend its estimation techniques for WCET bounds towards more complex architectures. Results will be integrated by partners.

The group at Bologna will continue to cooperate with STM (now an affiliated partner) on problems of industrial relevance. Furthermore, extended cooperation with Dortmund on memory modelling is on the horizon.

EPFL and ETH Zürich are planning to continue 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.

ETH Zürich will be working together with partners such as Bologna, Saarbrücken and Dortmund on an informal basis, in particular on the design of timing-predictable systems.

The Technical University of Denmark will extend its work on resource aware design space exploration. Furthermore, the cooperation with University of Bologna on system modelling infrastructure will be extended to include power models for more accurate exploration of power consumption. Finally, a proposal for a joint project between Technical University of Denmark and University of Linköping on analysis and synthesis of low-power fault-tolerant embedded systems has been filed. The aim is to consider efficient inclusion of fault tolerance given the tight resource constraints.

6.4.4 Meetings Planned

Resource Aware Design Activity Meeting : April 16th, 2007, 9:30 – 12:30, Nice, France.

A meeting of the Execution platforms cluster will be held May 14th in Linköping.

7. Cluster: Control for Embedded Systems

Cluster Leader: Karl-Erik Arzen (Lund)

The following is a description of the activities and overall objectives for the period: September 2006 – February 2008. The next reporting period will cover September 2006 – August 2007.

Staff Mobility

Area of Collaboration	**Feedback scheduling of control systems**
Sending Institution	UPC (Pau Marti) – ART cluster
Receiving Institution	LUND (Anton Cervin) – Control cluster
Persons	Rosa Castane Selga
Technical Work	Develop new feedback scheduling strategies
Dates	September – December 2006
Approximate Costs	Nb people : 1 Travel : 300 €

Stay: 8,000 € (paid by LUND, salary and accomodation)

Long Range Impact on Integration Important in order to improve the integration

Published Work Rosa Castañé, Pau Martí, Manel Velasco, Anton Cervin, Dan Henriksson. Resource Management for Control Tasks Based on the Transient Dynamics of Closed-Loop Systems. In Proceedings of the 18th Euromicro Conference on Real-Time Systems, Dresden, Germany, July 2006.

Further Collaboration Planned Yes

Area of Collaboration	**Tools for embedded control systems**
Sending Institution	CTU, CZECH REPUBLIC – Control cluster
Receiving Institution	LUND (Karl-Erik Årzén) and KTH (Martin Törngren) – Control cluster
Persons	Zdenek Hanzalek
Technical Work	Specification of common interface and case studies for design tools TrueTime and Torsche. The scheduling algorithms for FPGAs have been integrated in the concept of the TrueTime tool.
Dates	15.9.2005 – 24.9.2005
Approximate Costs	Nb people : 1 Travel : 250 €

Stay: 1,250 €

Long Range Impact on Integration Important in order to improve the integration

Published Work Martin Törngren, Dan Henriksson, Karl-Erik Årzén, Anton Cervin, Zdenek Hanzalek. Tools Supporting the Co-Design of Control Systems and Their Real-Time Implementation; Current Status and Future Directions. In Proceedings of the 2006 IEEE Computer Aided Control Systems Design Symposium, October 2006.

Further Collaboration Planned Yes

Area of Collaboration	**Sensor Networks**
Sending Institution	CTU , CZECH REPUBLIC
Receiving Institution	KTH
Persons	Jiri Trdlicka, Ing.
Technical Work	Development of the network flow routing algorithm with real-time constraints. Previous work by Mikael Johansson on optimisation of multi-commodity flows was extended by real-time constraints and experiments have been carried out in Matlab.
Dates	31.5. 2006 -29.8 2006
Approximate Costs	Nb people : 1 Travel : 300 € Stay: 5,500 €
Long Range Impact on Integration	Important in order to improve the integration
Published Work	J.Trdlicka, M.Johansson, Z.Hanzalek, Network flow routing algorithm with real-time constraints, internal report to be submitted to appropriate conference
Further Collaboration Planned	Yes

Area of Collaboration	**Real-Time Control**
Sending Institution	UPVLC, Spain (Pedro Albertos) – Control cluster
Receiving Institution	LUND (Karl-Erik Årzén) – Control cluster
Persons	Pedro Garcia
Technical Work	Development of new delay compensaytion schemes
Dates	Mid June – Mid September 2006
Approximate Costs	Nb people : 1 Travel : 350 € Stay: 1,250 €
Long Range Impact on Integration	Important in order to improve the integration
Published Work	Pedro Garcia, Pedro Albertos, Tore Hägglund. Control of unstable non-minimum phase delayed systems. Accepted for publication in Journal of Process Control. José Luis Guzmán, Pedro García, Tore Hägglund, Sebastián Dormido, Pedro Albertos, Manuel Berenguel. "Interactive tool for analysis of time-delay systems with dead-time compensation" In 7th IFAC Symposium on Advances in Control Education, Madrid, Spain, June 2006.
Further Collaboration Planned	Yes

Area of Collaboration	**Embedded control systems**
Sending Institution	KTH
Receiving Institution	CTU
Persons	Martin Törngren and Bengt Eriksson
Technical Work	Graduate course on Embedded Control Systems
Dates	3 – 7 April 2006
Approximate Costs	Nb people : 2 Travel : 600 € Stay: 1200 €
Long Range Impact on Integration	Important in order to improve the integration
Published Work	No
Further Collaboration Planned	Yes

Area of Collaboration	**Embedded control systems**
Sending Institution	LUND
Receiving Institution	CTU
Persons	Karl-Erik Årzén and Anton Cervin
Technical Work	Graduate course on Embedded Control Systems
Dates	3 – 7 April 2006
Approximate Costs	Nb people : 2 Travel : 500 € Stay: 1200 €
Long Range Impact on Integration	Important in order to improve the integration
Published Work	No
Further Collaboration Planned	Yes

Area of Collaboration	**Embedded control systems**
Sending Institution	UPVLC
Receiving Institution	CTU
Persons	Pedro Albertos and Alfons Crespo
Technical Work	Graduate course on Embedded Control Systems
Dates	3 – 7 April 2006
Approximate Costs	Nb people : 2 Travel : 600 € Stay: 1200 €
Long Range Impact on Integration	Important in order to improve the integration
Published Work	No
Further Collaboration Planned	Yes

Area of Collaboration	**Embedded control systems**
Sending Institution	KTH
Receiving Institution	CEA
Persons	Martin Törngren
Technical Work	Joint project discussions with CEA and Volvo
Dates	May 2006
Approximate Costs	Nb people : 1 Travel : 300 € Stay: 150 €
Long Range Impact on Integration	Joint research projects are important in order to be able to implement the vision of the network.
Published Work	No
Further Collaboration Planned	Yes

Area of Collaboration	**Embedded control systems**
Sending Institution	KTH
Receiving Institution	LTH
Persons	Carl Johan Sjöstedt
Technical Work	PhD student exchange
Dates	26-28 April 2006 + 22-24 May 2006
Approximate Costs	Nb people : 1 Travel : 300 € Stay: 300 €
Long Range Impact on Integration	Important for the future collaboration
Published Work	No
Further Collaboration Planned	Yes

Area of Collaboration	**Embedded control systems**
Sending Institution	KTH
Receiving Institution	Volvo
Persons	Martin Törngren
Technical Work	Joint project discussions
Dates	Two separate short meetings during Spring 2006
Approximate Costs	Nb people : 1 Travel : 300 € Stay: 300 €
Long Range Impact on Integration	Important for the future collaboration
Published Work	No
Further Collaboration Planned	Yes

Past Meetings in Year 2

Cluster working meeting in connection with the CDC-ECC conference in Sevilla

- Date: Wed 14 December 2005
- Borrowed location at local university
- Organized by Pedro Albertos, UPVLC
- Objectives: To organize the internal work in the cluster during the year. To discuss research issues
- Partcipants: Karl-Erik Årzén and Anton Cervin (LUND), Karl-Henrik Johansson and Mikael Johansson (KTH), Pedro Albertos and Alfons Crespo (UPVLC)
- Results: The work for the coming year was planned. Responsibilities were assigned.

Cluster working meeting in connection with Graduate School on Embedded Control Systems

- Date: Tue 4 April 2006
- CTU, Prague
- Organized by Zdenek Hanzalek, CTU
- Objectives: To follow-up the internal work in the cluster during the year. To discuss research issues
- Partcipants: Karl-Erik Årzén and Anton Cervin (LUND), Martin Törngren and Bengt Eriksson (KTH), Pedro Albertos and Alfons Crespo (UPVLC), Zdenek Hanzalek (CTU)
- Results: The status of the different cluster activities were checked.

Network activity “Adaptive RT, HRT and Control” working meeting

- Date: Saturday 11 March 2006
- SSSA Pisa
- Organized by Giorgio Buttazzo (PISA) and Anton Cervin (LUND)
- Objectives: To discuss the joint research work within the activity between the ART and the Control cluster
- Participants: Alan Burns, Gerhard Fohler, Pau Marti, Giuseppe Lipari, Giorgio Buttazzo, Anton Cervin
- Results: Topics for joint research were identified.

Industrial Workshop: **Interaction between control and embedded electronics in automotive industry**

Organized jointly with the R-T Components cluster. Associated with the Beyond AUTOSAR meeting.

- http://www.artist-embedded.org/artist/-ARTIST2-Workshop-Beyond-AutoSar-.html
- Date: Thursday 23 March, 2006
- Innsbrück, Austria
- Organisers: Karl-Erik Årzén, Albert Benveniste and Werner Damm
- Objective: To identify the main issues for embedded control in the automotive industry sector, in particular with respect to timing and component models

- Participants: The meeting attracted around 40 partcipants. Invited presentations were held by Stefan Kowalevski, Karl-Erik Årzén and Carlos Canudas de Wit. More information about the names of the participants are available through Albert Benveniste
- Results: A joint publication documenting the results of the meeting are under preparation. The conclusions of the meeting are summarized in the activity reports "Industrial Workshops" and "Adaptive RT, HRT and Control"

Summer School: **ARTIST2 Graduate Course on Embedded Control Systems**

- http://www.artist-embedded.org/artist/-ARTIST2-Graduate-Course-on-.html
- Date: 3-7 April, 2006
- CTU, Prague, Czech Republic
- Organisers: Zdenek Hanzalek, CTU
- Objective: Annual cluster summer school/graduate course on embedded control systems
- Participants: The course was lectured by Zdenek Hanzalek, Karl-Erik Årzén, Anton Cervin, Martin Törngren, Bengt Eriksson, Perdo Albertos, Alfons Crespo, and Vladimir Havlena. The course had 42 participants.
- Results: A report summarizing the course has been generated. It is available from the cluster review material page http://www.md.kth.se/RTC/ARTIST2/publications.html

Summer School**: First European Laboratory on Real-Time and Control for Embedded Systems**

- URL: http://www.artist-embedded.org/FP6/ARTIST2Events/Events/RT-Control/
- Date: 10-14 July, 2006
- Pisa, Italy
- Organisers: Giorgio Buttazzo and Giuseppe Lipari (Pisa) and Karl-Erik Årzén and Anton Cervin (LUND)
- Objective: Laboratory/graduate course on real-time and control organized jointly between the ART and the Control cluster.
- Participants & Results: A report summarizing the course has been generated. It is available from the cluster review material page http://www.md.kth.se/RTC/ARTIST2/publications.html

Scandinavian ARTIST2 day

- URL: http://www.snart.org/
- Date: 21 August 2006
- Stockholm, Sweden
- Organisers: Anton Cervin (LUND) and Martin Törngren (KTH)
- Objective: Disseminate information about the activities within ARTIST2
- Participants: All clusters were represented except the ART cluster. ARTIST2 presenters were Karl-Erik Årzén, Kim Larsen, Bengt Jonsson, Martin Törngren, Jan Madsen and Björn Lisper. The number of participants were around 70.

Results: The awareness of ARTIST2 was increased among the audience. Several interesting issues relating to industry-academia cooperation were identified.

Meetings Planned in Year 3

Cluster Meeting in Connection with the CACSD Conference in Munich

- Date: 6 Oct 2006
- Objective: To discuss the upcoming review. To discuss the common framework for flexible control
- Participants: Karl-Erik Årzén, Anton Cervin, Zdenek Hanzalek, Alfons Crespo, Martin Törngren.

Invited ARTIST2 session on Co-Design Tools at the IEEE CACSD Conference in Munich

- Date: 4-6 Oct 2006
- Jointly organized session that presents the advances reached for the ARTIST2 do-design tools for embedded control

Cluster Meeting in connection with the Y2 Review meeting

- Date: 7 Nov 2006

International Workshop on Control for Embedded Systems

- Date: January-February 2007
- Urbana-Champaign, US
- Follow-up to the successful Lund workshop 2005.

Graduate School on Embedded Control Systems

- May 2007
- Lund, Sweden
- One week graduate course

Cluster Meeting in connection with Graduate School

- May 2007
- Lund, Sweden

Industrial workshop together with RT Components cluster

- Topic possibly Control in aerospace systems
- During Spring 2007

Activity on Control for Real-Time Computing in Connection with FeBID

- Munich, May 2007

It is likely that there will be further meetings during Year 3.

7.1 Platform: Design Tools for Embedded Control

7.1.1 Year 1 Achievements: Sept 2004 – August 2005

A survey on co-design tools for modelling and design of real-time control systems has been completed. In addition to this a state of the art survey on approaches for model/tool integration and model management has been initiated.

The individual tools have been further developed. For example, the TrueTime tool from LUND has been extended with support for wireless network blocks, battery-powered devices, and local clocks with drift and offset. The TORSCHE tool from CTU and the tools from UPVLC have also been further developed.

The tools have been promoted. For example, course and training material has been developed for TrueTime and a tutorial on TrueTime was given at the IFAC World Congress, Prague, July 3.

7.1.2 Year 2 Achievements: Sept 2005 – August 2006

A detailed description of these achievements is provided in the activity's 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 http://www.elet.polimi.it/conferences/cca06/CACSD_home.htm
- **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.
- **Achievement: Tool Integration**
  An example of how the to co-design tools TrueTime and Jitterbug from LUND can be combined has been developed. In [**Erreur ! Source du renvoi introuvable.**]

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.

- **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, that 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 2,000 times.
  Torsche: The development of a simulation and implementation support for DSP applications in TORSCHE has started. As far as for the input side of TORSCHE, we have designed a language, compatible subset of Matlab, suitable for description of DSP algorithms. The parser of this language, generating the graph of precedence relations from the language description, has been designed in BISON and FLEX. 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 simple illustration of this work will be presented in [**Erreur ! Source du renvoi introuvable.**]. 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.

7.1.3 Objectives and Work Planned: Sept 2006 – February 2008

The planned work for the coming 18 months includes the following parts

- Further development of partner individual tools
- Further work on model and tool integration including
  - Development of integration scenarios
  - Case studies involving integrating of tool functionalities developed by cluster partners
  - Case studies providing integration with UML tools
  - Case studies providing integration with tools for system safety analysis
- Further dissemination of results

To create a better cross-cluster understanding, and map of tools for embedded systems development, it is our opinion that a joint tool/platform meeting involving all the clusters should be organised within Artist2.

Future developments of Torsche (CTU)

Currently we are working on the DSP code generator, which transforms a schedule produced by TORSCHE either to the code to be simulated in TrueTime or to the Handel C to be implemented in FPGAs. The code generator is based on XSLT transformation. Fully automated version of the code generator will be finished during one year.

A new version (0.3) of TORSCHE will be released in October. The version will include XML support, new graph algorithms, interface to SAT solvers, examples of interconnection with TrueTime and response time analysis. Further we will work on graph editor and web based production of scheduling results in Gantt charts written in Perl and Metapost.

Future TrueTime Developments

A drawback with the current version of Truetime is that it is not possible to simulate production code directly. Instead the code for each task must be manually translated into the code-segment structure of TrueTime. During the year preliminary investigations have been made on how this problem can be solved. The solution that is currently discussed is based on the possibility to have multi-threaded code in Matlab S-functions (using MEX-files). It appears to be possible to derive a solution that does not require any manual code tranformation at all. This opens up intersting possibilities. For example, it would then be possible to start by simulating a controller in Simulink againat a Simulink model of the plant. When that performs according to specifications, an existing tool such as Real-Time Workshop could be used to generate C code for this controller. This C-code could then be simulated executing as a task in a TrueTime real-time kernel together with the other tasks, and the true timely behavior of the control loop could be investigated. Another issue that has been investigated is the possibility to have hierarchically structured code functions. This will be investigated during the coming year.

Future KTH tool developments

The design of advanced embedded control systems requires a systematic approach in handling their increasing complexity and in particular integration of the different system aspects supported by different modeling languages and tools. The work on the model and tool integration platform will continue. In particular emphasis will be placed on linking Matlab/Simulink with UML environments. While Matlab/Simulink has its emphasis on control and functional design, the UML environments enables a representation of overall system architecture that is compatible to an architecture description language currently under development. There are several parts of this work, including identifying the needs and usages of the various UML models, and the definition of suitable mappings (transformations) from Simulink to the selected UML representations. In developing prototype tools practical issues such as model and exchange formats, and tool APIs have to be considered.

Given the time and resources (depending on complementary funding/projects), the following work will also be performed in the coming period:

- Integrating a suitable UML tool with the integration platform previously developed.
- Providing facilities to integrate formal analysis tools for timing analysis and logical correctness (model checking) with Simulink/UML models.
- Extending the previously developed tool integration platform to handle not only development but also software production and maintenance for distributed embedded systems, e.g. allowing a system configuration to be defined, built and downloaded to a target, dealing with software allocation to the different nodes subject to established configuration rules and optimization criteria.

7.1.4 Meetings Planned

Meetings for Year 3 will be decided at the annual consortium meeting, Nov 7th in Paris (just before the review).

7.2 Cluster Integration: Control in real-time computing

7.2.1 Year 1 Achievements: Sept 2004 – August 2005

Since this a rather new research area it was decided that the main activity during the first year should be the creation of a research roadmap. This roadmap has been completed. The aim of the roadmap is to chart the area, provide a common platform for the coming work, and to identify the most important research directions.

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 value 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. • RTC 2005, a workshop on real-time control and control of real-time computing systems was organized in association with ECRTS 05 at Mallorca. An invited session on control over sensor networks and control of sensor network resources (co-organized with RUNES) has been accepted for the IEEE Conf on Decision and Control and the European Control Conference, Sevilla, Dec 2005.

A new feedback scheduling method was developed for control loops by Dan Henriksson and Anton Cervin. A paper will appear at the CDC-ECC'05 in Sevilla – LUND. 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.

7.2.2 Year 2 Achievements: Sept 2005 – August 2006

A detailed description of these achievements is provided in the activity's deliverable.

- **Achievement: Dissemination of Roadmap Material**
  The concludions from the roadmap developed during year1 were summarized into the conference paper [1] 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 [2] has also 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 followup 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 queing 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. In [4] the new model types derived for queing 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. The initial work has been reported in [7].
- **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. In [5], the inital states of the controlled plants are taken into account by the feedback scheduler by including the inital 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 includes a theoretical framework for self-regulating protocol design [14], as well as detailed resource control strategies for specific network technologies [15]. The KTH group has also worked on on-line error control adaptation in networked applications [8], feedback-based error-correction in feedback-based networks [9], stability of window-based queue control with applications to mobile terminal download, [10], models for network congestion control [11], and distributed consensus algorithms [12]

7.2.3 Objectives and Work Planned: Sept 2006 – February 2008

The exact topics that will be investigated are to a large extent decided by the forces outside the control of the network. However, our aim is to work on control-based models and methods for queuing systems with applications in server systems, control of multi-tier server systems, feedback scheduling of control systems, feedback-based and hierarchical resource reservation schemes for embedded systems, control of communication networs in different settings, e.g., congestion control, control of tramit power in wireless networks, control-based error coding, and optimization-based network protocol design.

The applicability of control based approaches to automotive embedded systems will be investigated in the context of the DYSCAS project (www.dyscas.org). Scenarios associated with dynamic configurations will be investigated. Services and an architecture handling the required dynamic configurations will be developed in the course of which control based approaches will be considered for this type of applications.

On a more general scale it is important to increase the visibility for this type of research. Currently there is a large industrial interest in the US, but so far it has been more modest in Europe. A good vehicle for achieving this goal is the next FeBID workshop in Munich in May 2007. Here it could be possible to make a dedicated ARTIST2 activity aimed at European industry.

7.2.4 Meetings Planned

Meetings for Year 3 will be decided at the annual consortium meeting, Nov 7th in Paris (just before the review).

7.3 Cluster Integration: Real-time techniques in control system implementations

7.3.1 Year 1 Achievements: Sept 2004 – August 2005

Since this a rather new research area it was decided that the main integration activity during the first year should be the creation of a research roadmap. The aim of the roadmap is to chart the area, provide a common platform for the coming work, and to identify the most important research directions. The roadmap consists of approx 60 pages.

Another important integration activity was the 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 value input. A separate research agenda for the work within Artist2 was written collectively as the output from the workshop.

A third important integration activity was the Valencia Graduate Course on Embedded Control Systems in April where all the cluster members lectured and the course material was developed jointly.

Additionally, a number of civilities have been performed. RTC 2005, a workshop on real-time control and control of real-time computing systems was organized in association with ECRTS 05 at Mallorca. An invited session on control over sensor networks and control of sensor network resources (co organized with RUNES) has been accepted for the IEEE Conf on Decision and Control and the European Control Conference, Sevilla, Dec 2005. An invited session about the research in the cluster was organized at the IFAC World Congress, Prague, July 8. The IFAC Summer School on Control, Computing and Communication, Prague, June 27 – July 1 was co-organized by the cluster. A special session on Model Driven Engineering at Euromicro, Porto, August 30 – September 3 was organized by the cluster. A number of quality publications have been produced by the members of the cluster during the year. For example, Årzen and Cervin are co-authors of the RTSS 25 year anniversary article “Real-Time Scheduling: A Historical Perspective” (has appeared in the Real-Time Systems journal). Several of the cluster members are also authors of chapters in the recently published “Handbook of Networked and Embedded Control Systems” (Birkhäuser), with Årzén in the editorial board.

7.3.2 Year 2 Achievements: Sept 2005 – August 2006

A detailed description of these achievements is provided in the activity’s deliverable.

- **Achievement: Dissemination of Roadmap Material**
  The dissemination of the Roadmap on “Real-Time Control Techniques Implementation” has been performed in several conference papers and courses. However, the complete roadmap has yet not been disseminated. The second edition of the Embedded Control Systems Graduate Course held in Prague April 3-7 2006, provided the opportunity to deliver the Roadmap to the community and to extract the main issues in form of lectures. Additionally, several papers have been presented in different workshops and conferences.
- **Achievement: Scheduling and control co-design technique 1: Jitter reduction models**.
  In order to reduce the jitter in control systems several activities have been carried out:

Probabilistic analysys of the response time of a control task. The response time probabilistic analysis is focused on calculate the response time distribution of a periodic task without simulating all over the hyperperiod (H). Some preliminary results have shown that the response time distribution in the interval [0,t] with t<H is very close (with a low error) to the response time of the task in [0,H]. This behaviour can be observed for any periodic system. However, the window [0,t] is different for every task set. The future work is focused on two ideas of how to find the parameter t:

- Using relationships between temporal parameters of tasks (C,D,P). We have developed a method to calculate t that achieves a response time distribution very close to the total distribution with an error less than 0.25% and a reduction factor (t/H) of 74%.
- Using statistical theory to calculate the size of the sample t that represents the population (H).

Deadline minimisation. The deadline minimisation is used to strongly reduce jitter of control tasks, in a real-time control application. Task periods are usually set by the system requirements, but deadlines and computation times can be modified in order to improve system performance. Sensitivity analysis in real-time systems is focused on changes in task computation times, using fixed priority analysis. The aim of this work is to provide a sensitivity analysis for task deadlines in the context of dynamic-priority, pre-emptive, uniprocessor scheduling. This work permits to obtain a deadline minimisation method that achieves the maximum reduction. As undertaken in other studies concerning computation times, we also define and calculate the critical scaling factor for task deadlines.

- **Achievement: Scheduling and control co-design technique 2: Evaluation of different controller task models.**
  A simulated system consisting of three independent plants with different initial parameters has been used to compare the performance when different methods to reduce the jitter are applied. The system is controlled by a computer with limited computational resources. So, a linear digital controller is designed for each plant. The three plants are implemented as real-time tasks such that the overall control performance is optimised. The methods compared are results of the partners previous proposals as:
  - STM: Typical task's model. Each task controls one pendulum.
  - CO_US: Lund model.
  - IMF: UPVLC model.
  - ICOFU: Hybrid system between CO_US and IMF. Integrated model proposed by Lund and UPV.
- **Achievement: Control kernel**
  The control kernel deals with the essential control activities to guarantee the safe behaviour of the complete system. For this purpose, the control software can be arranged in different layers. At the level of the OS, activities to closing the loop and driving the system to a safe position should be included. At the top level, the control system may include several on-line controller options as well as supervising and optimising activities.
  The work carried out has been focused in the definition of the parameters to define the platform support and the implementation of the control kernel.

- **Achievement: Operating system support for embedded systems**
  The development of specific services in the operating system for embedded control systems is one of the issues to be considered in this activity. These services include:
  - Specific scheduling policies related to the proposed task model to minimise the output jitter.
  - Control middleware which includes services to support the control kernel concept and functionalities.
  - Supervisor to support several execution environments or domains
  - Memory management in embedded systems with memory constraints.

  Two main results have been obtained:
  - XtratuM: It is a supervisor which permits to create different domains spatial and temporal isolated. Currently, a domain is based on control applications based on Partikle and the other one is Linux.
  - Partikle: It is a new real-time kernel which includes specific services for control systems. The kernel concept has been implemented as a Control middleware.
- **Achievement: Developments in Sporadic Event-based control**
  Normally, controllers are designed assuming equidistant (periodic) sampling. This simplifies the design process greatly, since the sampled plant description becomes a linear time-invariant (LTI) discrete-time system (assuming that the continuous plant was also LTI). However, other sampling schemes could be beneficial. From a computing or network point of view, it makes sense to only sample or control when something significant has occured in the system. In this work, we have investigated sporadic control of a first-order system, and compared the resulting performance and resource usage with ordinary periodic control and with aperiodic control which has been studied before by Lund. It is found that some performance can be gained even in the case where the sporadic controller is only allowed to sample more seldom than the periodic controller.
- **Achievement: Optimal on-line scheduling of multiple state feedback controllers**
  Digital controllers are usually designed as periodic tasks that regularly perform their sampling, computation, and actuation activities. In severely constrained systems, a better approach might be to only control one plant at a time. In this work, we have proposed a nonpreemptive on-line scheduling policy that uses the measured state of each plant when deciding which plant to control. Deriving the scheduling policy is very time-consuming but can be done off-line, using a technique called relaxed dynamic programming. We have also compared the nonpreemive on-line policy against common periodic schemes on a set of real laboratory processes. The results show that the new scheme can give large performance improvements while at the same time allowing the background tasks run when the need for control is small.
- **Achievement: Scheduling of control calculations on FPGAs**
  To facilitate the FPGA design process CTU works on scheduling algorithms using very universal model, where tasks are constrained by precedence delays and relative deadlines. The precedence relations are given by an oriented graph, where tasks are represented by nodes. Edges in the graph are related either to the minimum time or to the maximum time elapsed between start times of the tasks. The NP-hard problem of finding an optimal schedule satisfying the timing and resource constraints while minimizing makespan C_{max}, is being solved using several approaches. The first one is based on Integer Linear Programming, the second one is implemented as a

Branch and Bound algorithm, the third one on budget-like heuristic algorithm and the fourth one on EDF-like heuristic algorithm.

- **Achievement: Time-Delay compensation**
  In practical digital implementation of any controller, delays appear due to transport phenomena, computation of the control input, time-consuming information processing in measurement devices, etc. The area of control of delayed systems has attracted the attention of many researchers in the past few years because delays may be responsible for instabilities in closed-loop control systems. In order to cope with these delays, a number of algorithms have been reported.
  The algorithm proposed by UPVLC is a discrete-time controller based on state feedback using the prediction of the state. A convergence analysis shows that the state converges to the origin in spite of uncertainties in the knowledge of the plant parameters, the system delay and even variations of the sampling period. The proposed control scheme also has been satisfactory implemented to control the yaw displacement of a real four-rotor mini-helicopter. The experimental validation has been developed on an embedded system, MaRTE OS, which allows the implementation of minimum real-time systems according to standard POSIX.13 of the IEEE.

7.3.3 Objectives and Work Planned: Sept 2006 – February 2008

The nature of NoEs makes it difficult to give any hard guarantees with respect to which type of technical work that will be done during the next year. Some of the topics that we aim to pursue are studies of the fundamental trade-offs that exist between sampling rates, delays, and jitter in networked control, event-triggered feedback control, future development of the control kernel concept, server-based implementation methods for control systems, optimization-based scheduling, and the definition of a common framework for the interaction between controllers and the underlying OS-middleware-hardware layer.

An important issue that is common to all the cluster activities is the organization of a follow-up workshop to the Lund Workshop on Control for Embedded Systems. This is planned for the early spring 2007.

Another important item for this activity is the annual Graduate School on Embedded Control Systems. The coming year it will be hosted by LUND (May 2007). This time our aim is to have two parallel tracks during the first day of the course: one for students with a control background and one for students with a computer science background.

7.3.4 Meetings Planned

Meetings for Year 3 will be decided at the annual consortium meeting, Nov 7th in Paris (just before the review).

7.4 NoE Integration: Adaptive Real-Time, Hard Real-Time, and Control

7.4.1 Year 1 Achievements: Sept 2004 – August 2005

Mälardalen (Fohler) and LUND (Cervin) are working on the integration of the jitter margin and flexible scheduling. The jitter margin is an extension of the classical delay margin to time-varying delays. The jitter margin, J(L), is defined as the largest input-output latency jitter for which closed-loop stability is guaranteed for any time-varying latency $\Delta \in [L, L+J(L)]$, where L is the constant part of the input-output latency. The jitter margin is based on small-gain theory, but is not particularly conservative. The stability test is expressed in the frequency domain and a simple graphical interpretation involving the magnitude of the frequency curve exists.

The jitter margin can be used to derive hard deadlines that guarantee closed-loop stability, provided that the scheduling method employed can provide bounds on the worst-case and best-case response times of the controller tasks. The jitter margin can also be used when selecting network protocol for networked control loops.

In the context of the flexible scheduling framework, the jitter margin is used as a design tool for finding optimal static schedules for multiple concurrent control loops. The proposed design procedure can be outlined in three steps. The first step is to find a static schedule for the controllers that give acceptable control performance in terms of the apparent phase margin. The second step is to use nonlinear optimization techniques to optimize the schedule with regard to the control performance, again as measured by the apparent phase margin. The third step is to compute how much additional jitter can be tolerated in each task and use this information to allow sporadic tasks to execute. The tasks are then scheduled on-line using the slot-shifting technique.

Lund (Cervin) and Ericsson (Eker) are continuing their development of the control server model, with focus on distributed systems. The control server is a scheduling mechanism tailored to control tasks that combines three different ideas:

- *Reservation-based scheduling*. Each task is scheduled by a modified constant bandwidth server, where a dynamic server period is used.
- *Subtask scheduling*. A task may be divided in several segments that are scheduled as subtasks. Scheduling the two main parts of a control algorithm (Calculate and Update) as subtasks, the input-output latency of the controller can be reduced.
- *Time-triggered I/O*. Inputs can be read and outputs can be written at predefined points in time by the kernel, minimizing the jitter in the control actions.

The basic resource allocation model in the Control Server uses the concept of shares. To allocate a share x% of a resource to a component means that that component appears to be using its own private resource with x% of the original capacity. This is called ideal resource reservation. Note that the resource allocation is uniform over time. This means that, even if a component is not using a resource during a time interval, the resource will still be allocated.

In the single-CPU case, time-uniform resource allocation makes sense. In distributed systems, however, this would incur a delay of one period per node—something which is intolerable for many applications. For networked applications, the delay in a node can be shortened by increasing the resource usage. The downside is that quite large amounts of slack may result in the schedule.

Another complication in the distributed case is that the same resource may be used several times by the same end-to-end task. In a wireless networked control loop, for instance, it is likely that the same shared medium will be used both for the transmission of the measurement signal (from the sensor node to the control node) as for the control signal (from the control node to the actuator node). Hence, some time scheduling is inevitable to prevent further resource waste.

Pavia/Pisa (Buttazzo) and Lund (Cervin) have collaborated on adding support for the control server model in the SHARK kernel with the aim to use SHARK as a shared platform for implementing control applications.

Collaboration has started involving Lund, UPC, and Mälardalen. Rosa Castane Selga is a PhD student candidate from UPC that is spending August – December 2005 in Lund working on feedback scheduling methods for control. At Lund a new feedback scheduling strategy for multiple control tasks has been developed that uses feedback from the plant states to distribute the computing resources optimally among the tasks. Linear-quadratic controllers are analyzed, and expressions relating the expected cost to the sampling period and the plant state are derived and used for on-line sample-rate adjustments. In the case of minimum-variance control of multiple integrator processes, an exact expression for the optimal sampling periods can be obtained. For the general case, an on-line optimization procedure is used.

In the collaboration this approach is extended in several directions. One extension is the combination of this feedback scheduling method with the control server model in order to achieve deterministic input output latencies in the control loop. Another extension is to use a cost function that is based on the integrated absolute error instead of quadratic cost functions.

Several integration activities were performed involving CTU, UPVLC and SSSA (affiliated to Pavia) related to the OCERA project. 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.

In order to add flexibility to the real-time applications UPVLC has developed a nano-kernel called Xtratum. Xtratum is a thin layer of software that provides a simple and convenient API to access interrupt mechanisms and timer devices. Xtratum permits the execution of environments/applications under spatial and temporal isolation. Xtratum has been developed under the OCERA project.

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 autogyro (data acquisition and visualization), and a robotic arm demonstrator (servo-control).

In addition, the ART, Control, and HRT cluster members were strongly involved in mutual workshops and seminars.

7.4.2 Year 2 Achievements: Sept 2005 – August 2006

A detailed description of these achievements is provided in the activity's deliverable.

- **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 Innsbrück, March 23. It was co-locacted 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.
- **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 preemption 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, we 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 [3].
  - Lund (Cervin) and Pisa (Bini) are working 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, we 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 inital states of the controlled plants are taken into account by the feedback scheduler by

including the inital 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, 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. 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 Kasierslautern (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 devloped 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 the Master Thesis [2] 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 [4]. UPC has also added new features to Shark to allow easy implementation of feedback scheduling [5].
- A strong research connection is currently being established between CTU and UPCLC in the Control cluster and UCantabria, Pisa, UPC, 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 Firrst European Laboratory on Real-Time and Control for Embedded Systems was organized in Pisa, Italy, July 10-14, 2006. The number of participants were 40.
  http://www.artist-embedded.org/FP6/ARTIST2Events/Events/RT-Control/

7.4.3 Objectives and Work Planned: Sept 2006 – February 2008

The objective of the Artist2 network integration activity **Adaptive Real-time, HRT and Control** is to integrate the research performed within the clusters on Adaptive Real-Time System, RT-Components, and Control for Embedded systems on different computational models for embedded control systems and on the use of control techniques to provide adaptivity and flexibility in embedded systems. Each of the clusters have matching internal cluster activities, e.g., in the Control for Embedded Systems the corresponding cluster activities are Real-time techniques in control system implementations and Control in real-time computing. The activities within the cluster can be characterized as follows. There are strong joint research activities between the ART cluster and the Control cluster. These activities will be continued also during the next 18 months. The interaction with the RT-Components cluster is mainly performed through jointly organized workshops with industry. This approach will be continued also during the next 18 months.

The research problems to be tackled during the next 18 months involves both the use of control-techniques in resource scheduling for embedded systems and scheduling techniques and computational models for embedded control applications. These two lines are also combined in the form of feedback-based scheduling of embedded control systems. The explicit research problems that will be pursued is to a large extent dependent on the particular goals of the research projects that provide the direct funding for the activities. It is, however, quite clear that there will be work on feedback-scheduling of control tasks, event-based control, further use of SHARK and TrueTime as experimental and simulation platforms for the joint work, and mechanisms for handling overruns within this activity.

The aim for the industrial workshop is to organize a follow-up workshop on embedded control issues within a particular industrial branch. It is likely that the workshop will be co-organized with the Industrial Workshop network activity and that the particular branch will be the aerospace industry. A likely point in time for this workshop will be during the Spring 2007. A report summarizing the conclusions of the workshop will be produced.

During Early Spring 2007 a followup workshop to the Lund workshop on Control for Embedded Systems will be held. This time the workshop will be given at University of Illinois, Urbana-Champage with Tarek Abdelzaher and Lui Sha as the hosts. The format will be the same as last year, i.e. mainly invited participants and a lot of room for discussions. Tentative persons to be invited from outside the ARTIST2 community are John Doyle, Richard Murray, Mark Spong, Tariq Basar, Prasad Kumar, Williams Sanders, Klara Nahrstedt, Karl Åström, Dawn Tilbury and Joe Hellerstein. A report summarizing the conclusions of the workshop will be generated.

7.4.4 Meetings Planned

Meetings for Year 3 will be decided at the annual consortium meeting, Nov 7th in Paris (just before the review).

8. Cluster: Testing and Verification

Cluster Leader: Kim Larsen (Aalborg)

The following is a description of the activities and overall objectives for the period: September 2006 – February 2008. The next reporting period will cover September 2006 – August 2007.

Staff Mobility

Laurent Doyen, a PhD student of J-F Raskin at Bruxelles, will join Tom Hanzinger's team at EPFL in October 2006 as a postdoc.

8.1 Platform: T&V Platform for Embedded Systems

8.1.1 Year 1 Achievements: Sept 2004 – August 2005

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 accessible via a common web portal (the Yahooda web-page maintained by Brno).

8.1.2 Year 2 Achievements: Sept 2005 – August 2006

A detailed description of these achievements is provided in the activity's deliverable.

- **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.

8.1.3 Objectives and Work Planned: Sept 2006 – February 2008

Based on the partner's further development of existing and new tools for quantitative testing and verification, the platform activity will focus on the following issues for the next 18 months:

- The work on tool evaluation through industrial case studies will be continued and reported in the web repository on a regular basis. Also, links to stable and mature versions of the tools will be provided and updated for download.
- As mentioned above, tackling the problems of mutual exploitation of European Grid resources for model checking and establishing a common web interface will be postponed for the time being. However, the established link to European Grid projects on verification will be maintained through regular meetings in order to pursue the overall vision of a powerful computing facility.
- Within ARTIST2, the challenge for establishing high performance resources will be pursued by exploiting resources that are immediately available, like e.g. the NorduGrid facility, which has two clusters in Aalborg that may be applied for experiments. In particular, the distributed version of Uppaal and the Devine tool will be made available on the 50-node PC cluster, and experiments will be made for exploiting the 52 Gbyte shared memory facility for analysing large models by single-CPU tools.

As for the individual tools and algorithms, the following will be worked on:

- Neijmegen will continue their work on the symbolic test generation tool TorXagit, which includes results on how to test transition system swith data.
- Brno will address the problem of how to extend of distributed verification methods to an inter-cluster setting with the aim to effectivelly make use large networks of heterogenous computers.
- IRISA will investigate how to do symbolic test selection using coverage criteria for automata extended with variables.
- Aalborg and Uppsala wil continue their further development of the Uppaal tool with focus on (among other subjects) test generation and dissemination through industrial case studies.

8.1.4 Meetings Planned

Meetings for Year 3 will be decided at the annual consortium meeting, Nov 7th in Paris (just before the review).

8.2 Cluster Integration: Quantitative Testing and Verification

8.2.1 Year 1 Achievements: Sept 2004 – August 2005

In the first period a number of achievements within quantitative aspects of testing and verification have been made. These include cost guided searching techniques, testing theories and their implementations in tools like Uppaal, IF and Torx. Also, work on testing of infinite state systems has been made by INRIA. Finally a number of industrial case studies have been undertaken.

Considerable effort has been made on dissemination of results: The partners have been very active as invited speakers on Quantitative Testing and Verification at a number of conferences, and they have co-arranged a Dagstuhl meeting on testing (September 2004) and a summer school on 'Modelling and Components, Testing and Verification, Static Analysis' (September 2005).Year 2 Achievements: Sept 2005 – August 2006

A detailed description of these achievements is provided in the activity's deliverable.

- **UPPAAL 4.0 Real Time Verification**
  UPPAAL 4.0 is the result of over two and half years of development and contains many neew 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. In UPPAAL 4.0 the modelling language is extended with user defined functions. These are fully integrated into the modelling language, and can access and modify all state variables. The syntax follows the style of C/C++/Java, and mmost control-flow constructs of C are supported. The modelling language is also extended in order that priorities and channels may be specified and dealt with during analysis. On the performance side, 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. Main team Aalborg University.
- **Timed channels systems**
  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 propose Communicating Timed Automata (CTA) to model such systems. The goal was to study the borderline between decidable and undecidable classes of channel systems in the timed setting. Our technical results include proof of decidability in the setting of one channel (equivalent to one-counter machines) and proof of undecidability in the setting of two or more channels. It is noted that in the untimed setting, these systems are no more expressive than finite state machines. This shows that the capability of synchronizing on time makes it substantially more difficult to verify channel systems. Main team involved: Uppsala University.

- **Test-based learning of timed behaviour**
  In this work, we extend previous work by considering the full class of event-recording automata, while still avoiding to base it on the (usually prohibitively large) region graph. Our construction deviates from previous work on inference of automata in that it first constructs a so called timed decision tree from observations of system behavior, which is thereafter folded into an automaton. Main team involved: Uppsala University.

- **Verification and conformance testing for reactive system.**
  The work studies the combination of verification and conformance testing for the formal validation of reactive systems. In particular focus has been on verification and selection of test cases that may detect both non conformance and violation of properties. Main team involved: IRISA.

- **Symbolic test selection for extended automata using abstract interpretation**. We continue this work line by improving test selection in our toolset STG. Main team involved: IRISA.

- **Symbolic Determinisation of Extended Automata.**
  In this work, we define a symbolic determinisation procedure for automata extended with symbolic data variables, which has applications in verification, testing, and diagnosis of infinite-state systems. Main team involved: IRISA.

- **Off-line test generation for real-time systems**
  We present experiences from 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 focus on testing the software implementing the session (WSP) and transaction (WTP) layers of the WAP protocol. These layers, and their surrounding environment, are described as a network of timed automata. To model the many sequence numbers (from a large domain) used in the protocol, we introduce an abstraction technique. We believe the suggested abstraction technique will prove useful to model other similar protocols with sequence numbers, in particular in the context of model-based testing. A complete test bed is presented, which includes generation and execution of test cases. It takes as input a model and a coverage criterion expressed as an observer, and returns a verdict for each test case. The test bed includes existing tools from Ericsson for test-case execution. To generate test suites, we use UPPAAL Cover— a new test-case generation tool based on the real-time model-checker UPPAAL. Main team involved: Uppsala University and Aalborg University.

- **Testing of programs with floating point numbers.**
  Conformance testing of a program with floating point numbers with respect to its specification with real numbers. Main team involved: IRISA.

- **Black-box testing of cryptographic protocols.**
  Compositional approach for checking secrecy and authenticity properties of cryptographic protocols integrating ideas from verification, conformance testing, and learning, applied to biometric passports. Main team(s) involved: IRISA.

- **Verification of Communication Protocols using Abstract Interpretation of FIFO queues**.
  This work proposes a new approach to the verification of infinite states communicating processes based on an approximate analysis of channel contents by regular languages. Main team(s) involved: IRISA.

- **Supervisory Control of Infinite Symbolic Systems using Abstract Interpretation**. This work investigates the control of safety properties on infinite systems, modelled by transition system with data variables. Main team(s) involved: IRISA and on similar topics VERIMAG.
- **Analysis of Priced (Weighted) Timed Automata**
  Timed automata are a well-established formalism for the modeling and analysis of timed systems. Recently a very useful extension of timed automata has been proposed: priced (or weighted) timed automata. Priced timed automata are natural models for embedded systems where, often, resources consumptions have to be modeled. Priced timed automata extend classical timed automata with a cost function. Timed automata and priced timed automata are models for closed systems, where every transition is controlled. If we want to distinguish between actions of a controller and actions of an environment we have to consider timed games on those formalisms.
- **UPPAAL Cora: Optimal Scheduling and Planning**
  UPPAAL Cora is a branch of UPPAAL which allows for efficient analysis of cost-optimal reachability of priced timed automata. The original algorithm used a symbolic A* algorithm using so-called priced zones as main datastructure. It has been shown how the simple structure of the linear programs encountered during this symbolic A* algorithm can be exploited in order to substantially improve the performance of the current algorithm. The idea is rooted in duality of linear programs and we show that each encountered linear program can be reduced to the dual problem of an instance of the min-cost flow problem. Experimental results show a 70-80 percent performance gain. A framework for providing priced timed automata models scheduling problems is given as well as experimental results illustrating the potential competitiveness of our approach compared to existing approaches such as mixed integer linear programming. Main team involved Aalborg University.
- **Robustness issues for timed and hybrid automata**
  We have introduced a parametric semantics for timed controllers called the ASAP semantics. This semantics is a relaxation of the usual ASAP (ASAP stands for ``as soon as possible'') semantics (also called the maximal progress semantics) which is a mathematical idealization that can not be implemented by any physical device no matter how fast it is. On the contrary, any correct Almost ASAP controller can be implemented by a program on sufficiently fast hardware. We have studied the properties of this semantics and show how it can be analyzed using the tool HyTech. Main team CVF, and triggered several cooperations and future work with LSV-Cachan.
- **Analysis of O-minimal Hybrid Systems**
  Recently, the control of hybrid systems has appeared as a new interesting and active field of research, and many results have already been obtained. O-minimal hybrid systems have been first proposed in as an interesting class of systems. They have very rich continuous dynamics, but limited discrete steps.

- **Refinement of abstraction for affine hybrid automata**
  In this research, we have shown how to automatically construct and refine rectangular abstractions of systems of linear differential equations. From a hybrid automaton whose dynamics are given by a system of linear differential equations, our method computes automatically a sequence of rectangular hybrid automata that are increasingly precise over-approximations of the original hybrid automaton. We have proved an optimality criterion for successive refinements. We also have shown that this method can take into account a safety property to be verified, refining only relevant parts of the state space. The practicability of the method is illustrated on a benchmark case study. Main team(s) involved: EPFL and CVF.

- **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 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. Main team(s) involved: CVF.

- **UPPAAL Tiga: efficient synthesis of winning strategies for timed games**
  We have proposed a first efficient on-the-fly algorithm for solving games based on timed game automata with respect to reachability and safety properties. The algorithm we propose is a symbolic extension of the on-the-fly algorithm suggested by Liu & Smolka [LS98] for linear-time model-checking of finite-state systems. Being on-the-fly, the symbolic algorithm may terminate long before having explored the entire state-space. Also the individual steps of the algorithm are carried out efficiently by the use of so-called zones as the underlying data structure. Various optimizations of the basic symbolic algorithm are proposed as well as methods for obtaining time-optimal winning strategies (for reachability games). Extensive evaluation of an experimental implementation of the algorithm yields very encouraging performance results. On-going research include compact representation of winning strategies using symbolic datastructures such as BDDs and CDDs as well as their translation to executable control programs. Main team(s) involved: Aalborg with Nantes and input from LSV-Cachan.

- **Synthesis with incomplete information**
  In this research, we have proposed a fixed point theory to solve games of imperfect information. The fixed point theory is defined on the lattice of anti-chains of sets of states. Contrary to the classical solution proposed by Reif, our new solution does not involve determinization. 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. 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.

- **Algorithms for the verification of infinite state systems**
  In this research, we propose an abstract interpretation based approach to solve the coverability problem of well-structured transition systems. Our approach distinguishes from other attempts in that (1) we solve this problem for the whole class of well-structured transition systems using a forward algorithm. So, our algorithm has to deal with possibly infinite downward closed sets. (2) Whereas other approaches have a non generic representation for downward closed sets of states, which turns out to be hard to devise in practice, we introduce a generic representation requiring no additional effort of implementation. Main team(s) involved: CVF and LIAFA.

- **Rectangular abstractions of hybrid automata**
  We showed how to automatically construct and refine rectangular abstractions of systems of linear differential equations. From a hybrid automaton whose dynamics are given by a system of linear differential equations, our method computes automatically a sequence of rectangular hybrid automata that are increasingly precise over-approximations of the original hybrid automaton. We proved an optimality criterion for successive refinements. We also showed that this method can take into account a safety property to be verified, refining only relevant parts of the state space. The practicability of the method was illustrated on a benchmark case study. Main team(s) involved: EPFL.

- **Quantitative similarity between timed systems**
  We defined quantitative similarity functions between timed transition systems that measure the degree of closeness of two systems as a real, in contrast to the traditional boolean yes/no approach to timed simulation and language inclusion. Two systems are close if for each timed trace of one system, there exists a corresponding timed trace in the other system with the same sequence of events and closely corresponding event timings. We showed that timed CTL is robust with respect to our quantitative version of bisimilarity, in particular, if a system satisfies a formula, then every close system satisfies a close formula. We also defined a discounted version of CTL over timed systems, which assigns to every CTL formula a real value that is obtained by discounting real time. We proved the robustness of discounted CTL by establishing that close states in the bisimilarity metric have close values for all discounted CTL formulas. Main team(s) involved: EPFL.

- **Logics for real-time games**
  We added freeze quantifiers to the game logic ATL in order to specify real-time objectives for games played on timed structures. We defined the semantics of the resulting logic TATL by restricting the players to physically meaningful strategies, which do not prevent time from diverging. We showed that TATL can be model checked over timed automaton games. We also specified timed optimization problems for physically meaningful strategies, and we showed that for timed automaton games, the optimal answers can be approximated to within any degree of precision. Main team(s) involved: EPFL.

- **Symbolic testing**
  Symbolic testing aims at the integration of action-based testing (or control-flow testing, or state-based testing) and data testing. This means that actions in a state-based model can be equipped with data-parameters, and that the sequences of allowed actions can be determined by predicates over these data parameters. Currently these two approaches to testing are not integrated

- **Action refinement**
  The testing theory that was developed for testing communication protocols is message based: the test events are the sending or receiving of a message. Testing of component-based software systems is object based: the test events are method invocations, or, more precisely, method calls and method returns. This difference in granularity, or atomicity of test events hampers the application of communication protocol testing methods to component-based testing.

- **A framework for test coverage semantics**
  A framework to express coverage measures that express how well a test suite covers such a specification.
  Since testing is inherently incomplete, test selection has vital importance. Coverage measures evaluate the quality of a test suite and help the tester select test cases with maximal impact at minimum cost. Existing coverage criteria for test suites are usually defined in terms of syntactic characteristics of the implementation under test or its specification. Typical black-box coverage metrics are state and transition coverage of the specification. White-box testing often considers statement, condition and path coverage. A disadvantage of this syntactic approach is that different coverage figures are assigned to systems that are behaviorally equivalent, but syntactically different. Moreover, those coverage metrics do not take into account that certain failures are more severe than others, and that more testing effort should be devoted to uncover

- **A testing theory for probabilistic processes**
  A first step in developing statistical testing techniques for systems with nondeterministic behavior. We introduce a notion of finite testing, based on statistical hypothesis tests, via a variant of the well-known trace machine. Under this scenario, two processes are deemed observationally equivalent if they cannot be distinguished by any finite test. We consider processes modeled as image finite probabilistic automata and prove that our notion of observational equivalence coincides with the trace distribution equivalence proposed by Segala. Along the way, we give an explicit characterization of the set of probabilistic executions of an arbitrary probabilistic

- **On-line testing of real-time systems**
  The online testing tool Uppaal-TRON 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. Future work includes tight integration with the UPPAAL graphical user interface. Main team(s) involved: Aalborg and related work at VERIMAG.

- **Quantitative Compositional Reasoning.**
  A framework for compositional reasoning about qualitative system properties.
  Compositional reasoning about qualitative system properties. We present a compositional theory of system verification, where specifications assign real-numbered costs to systems. These costs can express a wide variety of quantitative system properties, such as resource consumption, price, or a measure of how well a system satisfies its specification. The theory supports the composition of systems and specifications, and the hiding of variables. Boolean refinement relations are replaced by real-numbered distances between descriptions of a system at different levels of detail. We show that the classical boolean rules for compositional reasoning have quantitative counterparts in our setting. While our general theory allows costs to be specified by arbitrary cost functions, we also consider a class of linear cost functions, which give rise to an instance of our framework where all operations are computable in polynomial time. Main team(s) involved: Twente.

- **Checking robustness of timed automata.**
  An algorithm for robustness checking based on zones.
  We propose 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. This problem is known to be decidable with an algorithm based on detecting strongly connected components on the region graph, which, for complexity reasons, is not effective in practice. Our symbolic algorithm is based on the standard algorithm for symbolic reachability analysis using zones to represent symbolic states and can then be easily integrated within tools for the verification of timed automata models. It relies on the computation of the stable zone of each cycle in a timed automaton. The stable zone is the largest set of states that can reach and be reached from itself through the cycle. To compute the robust reachable set, each stable zone that intersects the set of explored states has to be added to the set of states to be explored. Main team(s) involved: Twente.

- **A Continuous-Time Bayesian Network Reliability Modeling, and Analysis Framework.**
  A framework for dynamic systems reliability modelling and analysis using continuous-time Bayesian networks.
  We present a continuous-time Bayesian network (CTBN) framework for dynamic systems reliability modeling and analysis. Dynamic systems exhibit complex behaviors and interactions between their components; where not only the combination of failure events matters, but so does the sequence ordering of the failures. Similar to dynamic fault trees, the CTBN framework defines a set of basic BN constructs that capture well-defined system components behaviors and interactions. Combining, in a structured way, the various basic Bayesian network constructs enables the user to construct, in a modular and hierarchical fashion, the system model. Within the CTBN framework, one can perform various analyses, including reliability, sensitivity, and uncertainty analyses. All the analyses allow the user to obtain closed-form solutions. Main team(s) involved: Twente.

- **Synthesis and Stochastic Assessment of Cost-Optimal Schedules.**
  An alternative to the EPT approach to generate schedules that take the possible failures of resources into account.
  We present a novel approach to synthesize good schedules for a class of scheduling problems that is slightly more general than certain existing scheduling problems. The idea is to prime the schedule synthesizer with stochastic information more meaningful than performance factors with the objective to minimize the expected cost caused by storage or delay. The priming information is obtained by stochastic simulation of the system environment. The generated schedules are assessed again by simulation. The approach is demonstrated by means of a non-trivial scheduling problem from lacquer production. The experimental results show that our approach achieves in all considered scenarios better results than the extended processing times approach. Main team(s) involved: Twente.

- **Reachability in priced probabilistic timed automata.**
  An algorithm for cost-bounded probabilistic reachability problem.
  This work presents an algorithm for cost-bounded probabilistic reachability in timed automata extended with prices (on edges and locations) and discrete probabilistic branching. The algorithm determines whether the probability to reach a (set of) goal location(s) within a given price bound (and time bound) can exceed a threshold p in [0,1]. We prove that the algorithm is partially correct and show an example for which termination cannot be guaranteed. Main team(s) involved: Twente.

- **State identification problems for finite-state transducers**
  The problems of state identification have been well studied for models such as Mealy machines where inputs and outputs are synchronous, or at least have a one-to-one correspondence. Real-time models such as I/O timed automata, on the other hand, can be often abstracted by finite-state transducers, where inputs and outputs are asynchronous.
  We studied state-identification problems for such models and provided initial results. Main team(s) involved: VERIMAG..

- **Conformance testing for real-time systems**
  We studied different properties of the conformance relation tioco used in the model-based testing framework for real-time systems used in the tool TTG (timed test generator). In particular under which conditions the relation is compositional (i.e., A conforms to A' and B conforms to B' implies that the parallel composition of A and B conforms to the composition of A' and B'). We also compared tioco with other real-time testing conformance relations proposed in the literature. Main team(s) involved: VERIMAG and related work at Aalborg.

- **Decentralized observation problems**
  A fundamental observation problem is, given a model of the system to be observed and a specification of the property to be observed, to check whether the property is observable (i.e., the observer can resolve potential ambiguities due to partial observation capabilities) and if so to (automatically) synthesize an observer. We studied this problem in various decentralized settings (i.e., where there are more than on observers) and provided decidability and undecidability results. Main team(s) involved: VERIMAG.

- **Generating Path Conditions for Timed Systems**
  We provide 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. Our method can also be used for the automatic generation of test cases for unit testing. The current generalization of the calculation of path condition for the timed case turns out to be quite tricky, since not only the selected path contributes to the path condition, but also the timing constraints of alternative choices in the code. Main team(s) involved: VERIMAG.

- **Allen Linear (Interval) Temporal Logic**
  Translation to LTL and Monitor Synthesis: We show how Allen's logic can be translated to LTL and how to synthesize automatically monitors for specifications in this logic. Main team(s) involved: VERIMAG.

- **Product Lines**
  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. The environments show both the inputs that a system can receive and indicate which of the system's responses are relevant for the environment. The latter trait, called color-blindness, creates new possibilities for system transformations in the specialization process. We demonstrate the use of the framework by applying it to two classes of realistic design languages. Main team(s) involved: Aalborg.

- **Compositional Verification Using I/O-Automata**
  We propose a new look at one of the most fundamental types of behavioral interfaces: discrete time specifications of communication—directly related to the work of de Alfaro and Henzinger. Our framework is concerned with distributed non-blocking asynchronous systems in the style of

8.2.3 Objectives and Work Planned: Sept 2006 – February 2008

The long-term ambition of the Testing and Verification cluster is to improve current industrial practice by continuous dissemination and improvement of existing powerful testing and verification techniques. Within the Quantitative Testing and Verification activity our aim to provide modelling formalisms, methods and tools which will allow *quantitative* aspects to be dealt with and utilized for verification and performance analysis at early design stages as well as for systematic approaches to the testing phase.

The planned work includes continuation of metrics for testing coverage, abstraction methods and compositional methods allowing properties of a composite system to be inferred from those of its components.

Also, based on existing powerful (real-time) verification techniques the new research challenges identified within the second year will be continued in the next period. This includes work towards maturing and further development of important topics such as optimal scheduling, monitoring and fault diagnosis, controller synthesis, robustness and implementability of quantitative models and analysis of hybrid models, stochastic and timed models.

The theoretical work will be supplemented by experimental work on tool prototypes and case studies.

In somewhat more detail we expect to tackle the following problems during the next 18 months:

Verification:

- Systematic construction of verification models for embedded systems
- Implementation of robust model-checking algorithms for real-time systems.
- Development of efficient symbolic representations for arithmetic sets.
- Study of the properties of automata-based symbolic representations of sets of integer and real vectors (Real Vector Automata, RVA).
- Development of efficient methods for iterating transducers.
- UPPAAL with asynchronous communication
- Implementation of zone-based verification engine for probabilistic timed automata.

Testing:

- Establishing a relation between (ioco) testing theory and assume/guarantee frameworks
- Symbolic test selection using coverage criteria and incorporating a technique for test-data selection.
- Application of work on coverage metrics to realistic case studies
- Testing theory and test selection for recursive programs.
- Test-based modeling, i.e. a model is inferred from test observations.
- Continued development of the test generation tools UPPAAL Tron,TTG, and TorX.

Abstraction and approximate methods:

- Approximate methods for verification of timed systems, in particular systems with buffers for asynchronous communication and resource sharing.
- Methods for automatic abstraction refinement for hybrid and probabilistic systems

Compositionality:

- Compatability checking between timed interfaces.
- Compositional backwards reachability methods for timed systems.

Robustness and implementability;

- Identification of tractable, robust models of quantitative systems, possibly based on theory of continuity and discounting.
- Code synthesis from timed models to executable code.

Controller synthesis and optimal scheduling

- Generation of compact code from winning strategies for timed games using symbolic datastructures.
- Efficient algorithms for synthesis of winning strategies for timed games with incomplete information
- Continued development of UPPAAL Tiga
- Revisit the automata theoretic approach to model-checking in the light of the research done on synthesis for incomplete information.
- Further developments on synthesis of robust controllers (incomplete information)

Priced / Weighted Timed automata

- Efficient algorithms dealing with multi-priced models
- Efficient algorithms for optimal infinite schedules.
- Continued development of UPPAAL Cora
- Optimal stategies for priced timed game automata with two clocks.

8.2.4 Meetings Planned

Meetings for Year 3 will be decided at the annual consortium meeting, Nov 7th in Paris (just before the review).

8.3 Cluster Integration: Verification of Security Properties

8.3.1 Year 1 Achievements: Sept 2004 – August 2005

The main achievements during the period from Sept 2004 to August 2005 can be summarized as follows:

- We developed a proof that the Dolev-Yao model is a sound abstraction of the complexity theoretic model for protocols that combine several cryptographic primitives. This is a major result as: 1.) it provides a justification of the existing automatic verification methods for security protocols and 2.) It allows bridging a gap between the cryptology community, that has its own models and definitions of security properties and protocols as well as proof techniques of security, which are essentially complexity-theoretic, and the formal methods community that uses the Dolev-Yao model for which automatic verification methods and tools have been developed.
- We organized an international workshop dedicated to the verification of security protocols. The workshop, named Workshop on the link between formal and computational models, took place in Paris from 23-24 June 2005 and was very successful. We had approximately seventy participants and twenty-two presentations. The worldwide groups working on the subject were represented with an important participation from Industry.
- We developed a data base for security protocols and their properties that accessible vie the internet : the Security Protocols Open Repository at http://www.lsv.ens-cachan.fr/spore/
- We have developed a methodology with tool support for the certification of Smart card applications. More specifically, the Common Criteria (CC for short) are an international standard widely used in the Smart Card sector. The CC defines seven levels of certification for EAL 1 to EAL 7 with increasing demand on formal proofs and testing. While there is a number of applications certified up to the EAL5 very few certification exist for the EAL 6 and almost none for the EAL 7. While certifying an application at the highest levels can be a strong marketing argument, the cost of the development and the evaluation of a product that fulfils the CC requirements for these levels are too high. In particular, while there exists tool support and a methodology for the lower levels such support is missing for the EAL 6 and EAL 7. We developed a tool supported methodology for the EAL 6 and EAL 7. This work has been done in a collaboration involving an industrial tool editor Trusted Logic, a smart card applications developer Axalto, an evaluation body CEA-LETI and two research labs CEA-LIST and VERIMAG. A patent concerning the developed methodology is actually under study.
- There is a variety of specification methods for security protocols. As a first step toward integrating verification tools that use these methods, we developed a classification and relation of the different existing specification methods (multi-set rewriting and process algebra).
- We have studied the expressive power of a process calculus that allows one to express arbitrarily many runs of ping-pong protocol thanks to the presence of recursive definitions. We have established a number of decidability results that indicate the limitations of automatic verification even in this simple setting. Most prominently, we show that our process calculus is Turing-powerful.

- We have developed a general verification method for security protocols that can handle unbounded sessions, unbounded message size and unbounded fresh nonce creations.
- We have developed a sound and complete inference system for bounded-sessions cryptographic protocols (the messages size is still unbounded), method that has been extended to take into account protocols that can use timestamps.
- We have considered the problem of access control for the Calculus of Mobile Resources due to Godskesen, Hildebrandt, and Sassone. We establish a type system that lets us establish security policies for processes and show that our type system satisfies the usual requirements of type preservation under reduction and safety (i.e. that well-typed processes cannot misbehave.) Moreover, we present a sound type inference algorithm that will let us extract minimal security policies.
- We have uses of standard model-checkers for analysing various security protocols (e.g. use of muCRL, SPIN and CADP) and for addressing security treats based on real-time issues (using UPPAAL).

8.3.2 Year 2 Achievements: Sept 2005 – August 2006

A detailed description of these achievements is provided in the activity's deliverable.

- **A logic for constraint-based protocol analysis.**
  The technical achievement is the design of language for specifying security properties together with a new algorithm for checking them
  Description: The outcome is a new constraint-based *tool* for the verification of security properties which allows one to specify the properties to check using a linear temporal language. Main team involved: Univ. of Twente.
- **A non-monotonic language for the specification of Trust Management policies.**

  We propose RT-, a new trust management language.
  Description: The outcome is new language which adds a restricted form of negation to the standard RT language, thus admitting a controlled form of non-monotonicity**.** Main team involved: Univ. of Twente.
- **A state-dependent access control system.**
  Description: The outcome is a model of state dependent access control. This is useful in many applications like for example, patient health records and employee. We have developed a software tool for verifying access control systems, which can check systems against specifications of the capabilities of users. Main team(s) involved: Centre Fédéré de Verification (actually, the team of Namur). The implementation is at: http://www.cs.bham.ac.uk/~mdr/research/projects/05-AccessControl/rw-xacml-1_6.tar.gz
- **Relating two standard notions of secrecy**.
  Description: We initiate a systematic investigation of situations where reachability-based secrecy entails strong secrecy. We show that in the passive case, reachability-based secrecy actually implies equivalence-based secrecy for signatures, symmetric and asymmetric encryption provided that the primitives are probabilistic. For active adversaries in the case of symmetric encryption, we provide sufficient (and rather tight) conditions on the protocol for this implication to hold. Main team involved: Verimag.

- **The CL-Atse Protocol Analyser.**
  Description: We have implemented the first complete decision procedure for detecting attacks on cryptographic protocols (in the case of finite sessions) using a XOR operator. The tool is available at http://www.loria.fr/equipes/cassis/softwares/AtSe/. The system also outperforms the other ones on standard cases. Main teams involved: LORIA.
- **Design of a combination techniques for handling several equational intruder theories in protocol analysis.**
  Description: This technique allows a hierarchy between the operators and has been applied to exponentiation operator with exponents ranging in an abelian group. Main teams involved: LORIA.
- **On key cycles.** Recent results on interpreting symbolic security proofs in more accurate computational model rely on the assumption that no keys cycle can be produced during an execution of the protocol. Main teams involved: LORIA.
  Description: We have shown that deciding the existence of key-cycle for a bounded number of sessions is NP-complete. The procedure also applies to protocols with timestamps.
- **Sandboxing in a distributed pi-calculus.**
  Description: We developed an extension of Hennessy and Riley's Dpi calculus with digital signatures and sandboxing with an associated type system that handles authentication. See http://vbn.aau.dk/fbspretrieve/4528056/article.pdf. Participants: Hans Hüttel, Morten Kühnrich.
- **Recursion and replication in ping-pong protocols.**
  Description: Theorems that describe to which it extent it is possible to use automatic verification techniques for ping-pong protocols with recursion or replication. See http://www.brics.dk/~srba/files/HS:JAR:05.pdf. Participants: Hans Hüttel, Jíři Srba.
- **Preliminary integrated framework for security and trust management.**
  Description: We developed a preliminary integrated framework based on process algebras and suitable inference systems for the modelling of security protocols as well as of access control and trust/reputation management policies. Main teams involved: CNR-IT.
- **Synthesis of enforcing mechanisms for security policies**
  Description: We developed a framework for the automatic synthesis of enforcing mechanisms for security policies. In particular, we modelled as process algebra operators, the security automata of Schneider as well as the edit automata of Ligatti et al. Main teams involved: CNR-IT.
- **Verification of security properties of cryptographic Application Program Interfaces (API).**
  Description: We developed a formal specification of IBM's security API (Common Cryptographic Architecture) and a computed-aided proof of its security. Main team(s) involved: VERIMAG.

- **Computational soundness of the symbolic model for cryptographic primitives**
  Description: In the symbolic model, cryptographic primitives are considered as operations on abstract data type. This is not only the case for the protocol but also for the adversary trying to break the protocol. Cryptographic primitives are, however, modelled more accurately by randomized algorithms, and the security of a protocol is defined as the low probability that a probabilistic adversary with limited resources can break the protocol. Proving soundness of symbolic allows to benefit from the automated tools of the symbolic model on one hand and from the fact that the computational model is quite close to real implementations. We have proved the soundness of the symbolic model for protocols that use asymmetric and symmetric encryption, digital signature, hash functions and Diffie-Hellman exponentiation. Main team(s) involved: VERIMAG.
- **Development methodology with tool support that allows certification of Smart Card applications at the highest level EAL7 of Common Criteria.**
  Description: A computed-aided methodology for checking the formal conformance of applications with respect to security policy. We have extended the certification methodology to take into account new features both of applications (for example we can now handle more complex data structures) and of the security policy we want to check (data flow oriented properties in addition to trace-based security properties). The methodology is now being transferred to TrustedLogic and integrated in their tools. This transfer is financed by a French national project. Main team(s) involved: VERIMAG.
- **Certifying Cryptographic Protocols by Abstract Model-Checking and Proof Concretization**
  Description: The aim is to produce a proof of correctness independently from the tool and the abstractions used. First a proof of the abstract property is produced, and then it is automatically transformed into a proof of the concrete property and a set of proof obligations. Main team(s) involved: VERIMAG.
- **Specification language for cryptographic protocols.** We developed a specification language which makes it possible to separate the roles of a protocol from the scenario which defines how instances of the roles are created. In our system, roles are programs written in simple imperative programming language and are executed by (legitimate) protocol participants. Main team(s) involved: VERIMAG.
  Description: The outcome is a new specification language for cryptographic protocols which allows describing both protocols and the specific context in which they are used.
- **Formalization of protocols for electronic voting**
  Outcome: some protocols formalization including one from France Telecom. In a simplified model we get automatic proof of some security properties (fairness and eligibily)and by-hand proof of some other properties (receipt-freeness and coercion-resistance). Main teams involved: (France Telecom, LSV, INRIA, Univ. of Birmingham)

- **HERMES: A verification tool for cryptographic primitives**
  Description: HERMES is a tool for the automatic verification of cryptographic protocols. The initial version of HERMES implemented a general verification method based on abstraction, which can handle unbounded sessions, for protocols described using an Alice-Bob like specification language. The second version takes as input the new specification language mentioned above. The verification capabilities of HERMES have been extended with methods that handle specified scenarios (for example, unbounded but only iterative sessions, or composition between bounded and unbounded sessions, etc.).This second version allowed us to validate the protocol for electronic purse provided by France Telecom. The HERMES tool, versions 2, is available online at http://www-verimag.imag.fr/~Liana.Bozga/home/hermes.html

- **Symbolic Protocol Analysis in Presence of a Homomorphism Operator and Exclusive Or.**
  Description: We show that symbolic trace reachability for well-defined protocols is decidable in presence of the exclusive or theory in combination with the homomorphism axiom. These theories allow us to model basic properties of important cryptographic operators. Involved: LSV, LIF, Marseille.

- **A Survey of Algebraic Properties Used in Cryptographic Protocols.**
  Description: A great deal of cryptographic protocols relies on algebraic properties. We give a list of some relevant algebraic properties of cryptographic operators, and for each of them, we provide examples of protocols or attacks using these properties. We also give an overview of the existing methods in formal approaches for analyzing cryptographic protocols. Persons involved: V. Cortier (LORIA), S. Delaune (LSV) and P. Lafourcade (LSV).

- **Easy Intruder Deduction Problems with Homomorphisms.**
  Description: We present complexity results for the verification of security protocols. We are interested in theories such as Exclusive or and Abelian groups in combination with the homomorphism axiom. We show that the intruder deduction problem is in PTIME in both cases, improving EXPTIME existing results. Persons involved: S. Delaune (LSV).

- **Tree automata with equality constraints modulo equational theories.**
  Description: We present new classes of tree automata combining automata with equality test and automata modulo equational theories. This class has a good potential for application in software verification and is very useful, in the context of cryptographic protocol verification, to model the algebraic properties of the cryptographic primitives. Involved: LSV and LORIA.

- **Intruder Deduction for the Equational Theory of Exclusive-or with Commutative and Distributive Encryption**
  Description: The paper solves the intruder deduction problem (passive case) for a theory of Exclusive-or with commutative and distributive encryption. It is shown that this problem is in 2-EXP-Time and that even the binary case is EXP-SPACE-hard. Involved: LSV.

- **ACUNh: Unification and Disunification Using Automata Theory**
  Description: We propose an efficient decision procedure for the (dis)unification modulo the theory of the Exclusive-or with homomorphism. The algorithm follows an automata-theoretic approach. Involved: LSV and LIF, Marseille.

- **Guessing Attacks and the Computational Soundness of Static Equivalence.** *Description:* We give a computational justification of the use of a particular equational theory in the context of guessing attacks. Guessing attacks are formally modelled using static equivalence. Involved: LSV, LORIA, Microsoft research and UC Santa Cruz.
- **Coercion-Resistance and Receipt-Freeness in Electronic Voting.** *Description:* In the context of our efforts to formally study electronic voting protocols, we have studies prominent anonymity properties of election protocols; we have given formal definitions of privacy, receipt-freeness and coercion-resistance in the applied pi calculus. Involved: LSV, Univ. of Birmingham.

8.3.3 Objectives and Work Planned: Sept 2006 – February 2008

Our goal is to *broaden the horizon of the verification on security protocols* in such a way that it meets the requirements and the (future) expectations of industrial partners. As mentioned in section 1.5, this goal is made concrete in a threefold challenge:

1) The verification of more realistic protocols.
2) The verification of more realistic security properties.
3) Bridging the gap between the verification of security properties and trust management.

Concerning challenge (1), there are various problems we intend to tackle. First the verification of *group protocols:* nowadays, protocols often involve groups (whose size is not defined a priori) of participants. Typical examples of such protocols are group-key exchange protocols. Verifying such protocols is a major challenge, because the number of participants is a parameter. Moreover, new security properties emerge due to the fact that members can dynamically enter or leave a group. We wish to define a framework for defining and analysing such protocols and their related properties. Another problem we intend to tackle in challenge (1) is to lift the analysis of protocol properties to services properties: security protocols are often used in conjunction with other applications (e.g. access control) and may manipulate complex data (e.g. XML), in order to compose a service. We plan to address these service verification problems. Next to this topic, the problem of the verification of security properties should be broadened to include authentication protocols for mobile ad-hoc networks in the applied pi-calculus. Finally, we intend to develop a tool for the automatic verification of cryptographic APIs

Concerning challenge (2) Most automatically verifiable properties are reachability properties, such as secrecy and authentication. Anonymity properties and stronger versions of secrecy can be modelled elegantly using equivalence relations, such as observational equivalence in the applied pi calculus. We wish to define symbolic semantics of the applied pi calculus in terms of constraint systems and a corresponding symbolic observational equivalence relation. This should lead to new decidability and complexity results, as well as algorithms, for deciding this equivalence in the case of a bounded number of sessions and particular equational theories. A major advantage of equivalence based properties is that they are compositional. Within the same challenge, we are going to assess the usability of type inference algorithms for checking security properties: we intend to find and analyze a type inference algorithm for correspondence assertions in the spi-calculus and generalizing it to the applied pi-calculus. Further up in the goal line, we aim at implementing a tool that uses type inference for the analysis of cryptographic protocols.

Concerning challenge (3), we plan to define a complete and uniform framework for the specification of security protocols and trust management systems in complex, dynamic and open scenarios. The framework will be both supported by modelling and analysis tools as well as by effective implementations. A second problem that will be addressed in the coming months is the integration of rule-based and reputation based trust management systems.

The ongoing work on the individual tools will be continued. However, emphasis will also be made on evaluation of the tools through case studies in order to identify the most stable and mature versions with respect to integration.

8.3.4 Meetings Planned

Meetings for Year 3 will be decided at the annual consortium meeting, Nov 7th in Paris (just before the review).