ARTIST 2
Network of Excellence

IST-004527 ARTIST2:
Embedded Systems Design

Activity Progress Report for Year 1

JPRA-NoE Integration:

Semantic Framework for Hard Real-Time Design Flow

Clusters:

Hard Real-Time

Adaptive Real-Time

Control for Embedded Systems

Activity Leaders:

Albert Benveniste (INRIA) and Alberto Sangiovanni-Vincentelli (PARADES)

Large European systems industries must maintain their competitive position in the future. This requires improving substantially the entire OEM-supplier chain and the design methodology used to develop embedded systems. The methodology has to take into consideration that industry must completely revisit the way systems are decomposed into subsystems for to facilitate development by suppliers, and integration by the OEMs. This move will require that virtual design (i.e., design based on computer modelling and analysis) be performed systematically to discover errors at early stages of systems development. Using virtual engineering will require changing the technology, skill set, and support make-up of industry in a profound way. This course of action will require mastering heterogeneity in large design flows involving concurrent activities. Therefore, we need deep innovations in architectural abstractions capturing functional and non-functional features, in formal modelling (semantic based integration of heterogeneous system models with component models covering all non-functional constraints), and in formal multi-viewpoint analysis covering functional, timeliness, safety, and dependability requirements performed across all system design abstraction levels.

The objective of this action is to uncover the semantic issues involved with the use of heterogeneity at the functional and architectural level and point out an optimized course of action to answer them. We contend in fact, that many problems faced by industry in bringing new products to market can be traced to fundamental issues in the use of heterogeneous components that are answered at best in a heuristic and shallow fashion.

Table of Contents

1. Introduction

1.1 Activity Leader

Team Leader: Albert Benveniste (INRIA) and Alberto Sangiovanni-Vincentelli (PARADES)
Areas of their team's expertise:

- INRIA: synchronous languages, heterogeneous embedded systems;
- PARADES: platform based design, embedded systems and hardware.

1.2 Clusters

Hard Real-time
Adaptive Real-time
Execution Platforms

1.3 Policy Objective

European automotive and aeronautics industries are experiencing an exponential growth in functionality, with a drastic increase of innovations realized in software. The paradigm is shifting from an original “1 function = 1 ECU = 1 supplier” partitioning to distributed realizations of functions across multiple ECU’s involving multiple suppliers.

Several de-facto standards such as CAN, Flexray and the various OSEK extensions have found there way into series development. Model based development is increasingly gaining momentum, often involving automatic code generation. These processes are reaching substantial levels of maturity for single ECU implementations, including advanced capabilities for rapid prototyping and (hardware-in-the-loop) testing.

However, system-oriented design and virtual integration are supported only weakly, leading to late detection of integration problems. The OEM-supplier relation is mainly relying on textual initial requirements and well-established processes of delivering increasingly mature prototypes. Late requirement changes, incomplete initial requirements, or even inconsistent requirements are often leading to late design iterations or changes, with high-incurred costs.

Already today, privately funded key initiatives like Autosar demonstrate the commitment of this industrial sector to reducing costs. Way to this is by harmonizing platforms and decoupling functional architecture design from target platforms.

While these initiatives are demonstrating the industrial pull in the required direction, they are only making the first move. Challenges required to achieve the targeted growth rates, and not yet achieved, include:

1. The need for boosting re-uses across all design levels. This requires component models capturing the complete space of non-functional constraints (time, dependability, safety, scheduling, resource consumptions), as well as functional and protocol aspects in order to achieve drastic cost reductions.
2. The need for ensuring high quality and optimizing electronics despite the exponential growth in system complexity. This requires strong advances in enhanced virtual (sub)system integration and analysis, in order to reduce the number of deep iteration loops. Since subsystems are typically developed by multiple suppliers, this entails the need to integrate models from different modeling tools.

3. The need to optimize cross-supplier development processes. This requires early assessment of risks caused by late requirement changes, fast turn-around times in integrating resulting changes, and design space exploration across boundaries of the supplier chain.

To summarize, large European systems industries must maintain their competitive position in the future. This requires improving substantially the entire OEM-supplier chain and the design methodology used to develop embedded systems. The methodology has to take into consideration that industry must completely revisit the way systems are decomposed into subsystems for to facilitate development by suppliers, and integration by the OEMs. This move will require that virtual design (i.e., design based on computer modelling and analysis) be performed systematically to discover errors at early stages of systems development. Using virtual engineering will require changing the technology, skill set, and support make-up of industry in a profound way. This course of action will require mastering heterogeneity in large design flows involving concurrent activities. Therefore, we need deep innovations in architectural abstractions capturing functional and non-functional features, in formal modelling (semantic based integration of heterogeneous system models with component models covering all non-functional constraints), and in formal multi-viewpoint analysis covering functional, timeliness, safety, and dependability requirements performed across all system design abstraction levels.

The objective of this action is to uncover the semantic issues involved with the use of heterogeneity at the functional and architectural level and point out an optimized course of action to answer them. We contend in fact, that many problems faced by industry in bringing new products to market can be traced to fundamental issues in the use of heterogeneous components that are answered at best in a heuristic and shallow fashion.

1.4 Industrial Sectors

Addressing heterogeneity in embedded systems design is essential in the industrial sectors where large systems are built, involving multiple skills with different paradigms and tools. This includes the following sectors, with the first two sectors being the leading ones:

- Avionics (Event and Time-triggered systems are developed and effectively used at Airbus Industries);
- Automobile;
- Rail Transport;
- Energy Production.

2. Overview of the Activity

2.1 Artist Participants and roles

Team Leader: Alberto Sangiovanni-Vincentelli (PARADES)
Areas of his team's expertise: strong interaction with automotive, design software and semiconductor industry; expertise in design flows, tools and modelling methodologies with particular attention to Hard Real-Time; Platform-Based Design and Metropolis design framework for integration of design processes from OEMs to suppliers involving functional and non functional aspects.

Team Leader: Albert Benveniste (INRIA)
Areas of his team's expertise: synchronous languages and heterogeneous systems modelling and deployment.

Team Leader: Hermann Kopetz (TU Vienna)
Areas of his team's expertise: inventor of the TTA concept.

Team Leader: Werner Damm (OFFIS)
Areas of his team's expertise: embedded system modelling and validation, deep involvement in cooperation with the automotive industries.

Team Leader: Paul Caspi (Verimag)
Areas of his team's expertise: synchronous languages and heterogeneous systems modelling and deployment; tight cooperation with Airbus.

2.2 Affiliated partners and Roles

Team Leader: Jan Romberg (TU Munich)
Areas of his team's expertise: synchronous dataflow notations and tools, distributed architectures in automobile.

Team Leader: Luciano Lavagno (Politecnico di Torino)
Areas of his team's expertise: IC design and algorithms for synchronous and asynchronous design.

Team Leader: Stefan Kowalewski (Bosch)
Areas of his team's expertise: automotive industrial case study.

Team Leader: Jakob Axelsson (Volvo)
Areas of his team's expertise: automotive industrial case study.

Team Leader: Team Leader: Francois Pilarski (Airbus France)
Areas of his team's expertise: avionics industrial case study.

2.3 Starting date, and expected ending date

Starting date: September 1st, 2004

Expected ending date: September 1st, 2006

2.4 Baseline

Large European systems industries must maintain their competitive position in the future. This requires improving substantially the entire OEM-supplier chain and the design methodology used to develop embedded systems. The methodology has to take into consideration that industry must completely revisit the way systems are decomposed into subsystems for to facilitate development by suppliers, and integration by the OEMs. This move will require that virtual design (i.e., design based on computer modelling and analysis) be performed systematically to discover errors at early stages of systems development. Using virtual engineering will require changing the technology, skill set, and support make-up of industry in a profound way. This course of action will require mastering heterogeneity in large design flows involving concurrent activities. Therefore, we need deep innovations in architectural abstractions capturing functional and non-functional features, in formal modelling (semantic based integration of heterogeneous system models with component models covering all non-functional constraints), and in formal multi-viewpoint analysis covering functional, timeliness, safety, and dependability requirements performed across all system design abstraction levels.

The objective of this action is to uncover the semantic issues involved with the use of heterogeneity at the functional and architectural level and point out an optimized course of action to answer them. We contend in fact, that many problems faced by industry in bringing new products to market can be traced to fundamental issues in the use of heterogeneous components that are answered at best in a heuristic and shallow fashion.

2.5 Technical Description

Currently there is no engineering practice of correctly handling heterogeneity in embedded systems design. This is considered one of the important bottlenecks in design processes and smooth integration of partial designs or sub-systems, particularly when originating from different suppliers.

Attempts to address this exist in the academic community, including the work done around Ptolemy and Ptolemy II at Berkeley, the Metropolis project at Berkeley, Cadence and PARADES. The only other additional attempt that we know from industry is the RT-Builder tool by TNI-software.

Fundamentals supporting these approaches include:

- *Interface automata* types of approaches, where low level descriptions of interactions between systems obeying different Models of Communication and Computation (MoCC) is provided as a “semantic bus”.
- *Emerging deployment theories,* e.g., for analysing how a synchronous design can be deployed over a GALS (Globally Asynchronous Locally Synchronous) architecture.

This JPRA collects all key European teams in the area. It has taken advantage of tight interactions with other relevant clusters, mostly Components and Execution Platform.

The objective of this JPRA is to draw novel research directions in support of heterogeneity, in the form of fundamentals on modelling and design paradigms.

2.6 Organization of the report

The rest of the report is organized as follows.

Section 3 provides the activity report. The added value of having ARTIST2 with this activity is explained in this section. The funding of the reported activity is mainly ARTIST2, except that a few travels may have in some cases been supported by other sources.

Section 4 summarizes ongoing activity at the partners. This is report on ongoing research; it is not funded by ARTIST2 and is just provided for reference.

3. Activity Progress Report

3.1 Work achieved in the first 6 months

An important meeting was held in Rome and hosted by PARADES (January 2005), jointly with the JPRA-NoE on Merging ET with TT.

Aim of this meeting was to review approaches by the various partners. In addition, a few industrialists were invited, from GM-USA and BMW-Germany. The meeting was organized in the following way:

- Detailed minutes were recorded, discussed and approved at the meeting, and possibly enriched with additional material from some partners after the meeting. This resulted in very useful minutes that we attach to this report.
- Three long presentations were given by our industrial participants. These presentations were more focused on the topic of Merging ET with TT than on that of Semantic Platform. In these presentations, industrialists were expressing a number of concerns that we summarize below, and proposed a number of research directions, for the community.
- A number of long and detailed technical presentations were given by the academic participants, with slides provided. Some of the presentations gave rise to extended and hot discussions between participants.
- The minutes begin with an *executive summary* that collects the main findings from this meeting.

We collect here what we consider to be the major findings as collected in the above mentioned executive summary. ARTIST2-HRT and this JPRA must be credited for these findings.

3.1.1 Summary of research suggestions

The following trends have been identified by the automotive industry:

- Electronics is a significant component of vehicle, both in cost & complexity;
- It is growing at alarming rate (40% annually);
- Innovation is outpacing our ability to forecast.

GM competitive position is the following: GM is a high volume, low margin company historically; it has a most diverse portfolio. It faces an increasing competitive pressure with declining market share.

Therefore, GM has chosen the following technical strategies:

- Reuse ECUs when possible (all vehicles, all model years), with no artificial bounds, no dependency on forecast of future; this implies immunity to deployment platform differences.
- Separate logical feature content from physical deployment platform; use standard control infrastructure SW, and deploy OTS feature content to optimal physical architectures.

Control representation is provided that contributes to functional features. The latter must be deployed on physical ECUs; there, functional partitioning is a key issue. This partitioning may differ from vehicle to vehicle. Perform both logical and ECU reuse as much as possible;

From these considerations, the following research directions emerged:

- P. Caspi: some presentations related heterogeneity & scheduling. This blending seems to be an interesting line of research.
- P. Caspi: modeling and implementation; what are the different refinement relations between them?
- P. Caspi: MoCC, can we structure them? How many of them? Classification scheme of MoCCs? Why do they exist? Which ones should be distinguished? [to this end, there is a need to synchronize with clusters on components and executions platform]
- A. Sangiovanni-Vincentelli: comparison between environments that embed flows versus stand-alone flows, as paradigms to address embedded systems. Flexibility versus optimality in the tool space (of course both is best!). Metamodeling as a way of capturing heterogeneity in a formal way so that analysis and synthesis can be performed on sound basis.
- W. Damm: develop a new model of component that is rich enough to support both functional and non-functional characteristics, in a coherent and semantically sound way.

3.1.2 Future plans

It was decided that the next meeting should be common with other clusters, as indicated below.

3.2 Work achieved in months 6-12

A second meeting was held in Rennes and hosted by INRIA (June 2005), with other clusters participating: Components cluster and Execution platform cluster.

Again, this meeting was a joint meeting of the two JPRA-NoE on “Merging ET with TT” and “Semantic Platform”. This second meeting had its main focus on *Real-Time Components* and therefore was less addressing the issue of Merging ET with TT. We therefore mainly develop the outcomes of this meeting in this report.

3.2.1 Summary of research suggestions

- W. Damm (OFFIS): V-based development process is simple and familiar, but it indeed applies very concurrently. And this concurrency aspect is usually not considered seriously. How can one improve this? More precisely, how can we deal properly with the synchronizations under uncertainty that result from concurrency in the design process? Results today in deep iterations and costly recalls. How to maintain high degree of concurrency while mitigating the disturbances?
- W. Damm (OFFIS): develop research toward *rich components,* i.e., components that can support both functional and non-functional aspects in a coherent and compositional way, with semantically sound basis.
- J. Sifakis (Verimag, Components cluster): we need to develop a theory of *real-time architectures*, where some properties should be automatically preserved by components composition, e.g., deadlock freedom.

- L. Thiele (ETH, Execution Platforms cluster): dealing with *interfaces for real-time* raises novel issues and requires a new calculus. Fortunately, the area of QoS for networks provides us with a useful background to exploit.
- T. Henzinger (EPFL): the use of rigid logical time proved to be the basis for the success of synchronous languages. The same is needed for real-time. The idea is that start end termination times for tasks should be *both* fixed at compile time and not execution-dependent as it usually is for tasking systems.

We collect below the high-level approach to support research directions regarding components for real-time. This can be considered in part to be the outcome of ARTIST2-HRT cluster and JPRA on Semantic Platform. The figures below were collected and prepared for the EU-US day on components that was organized last July.

The following two figures show what the engineer would like to perform in a secure way. The context considered originates from automobile industry, but the arguments are valid in a more general setting, mutatis mutandis.

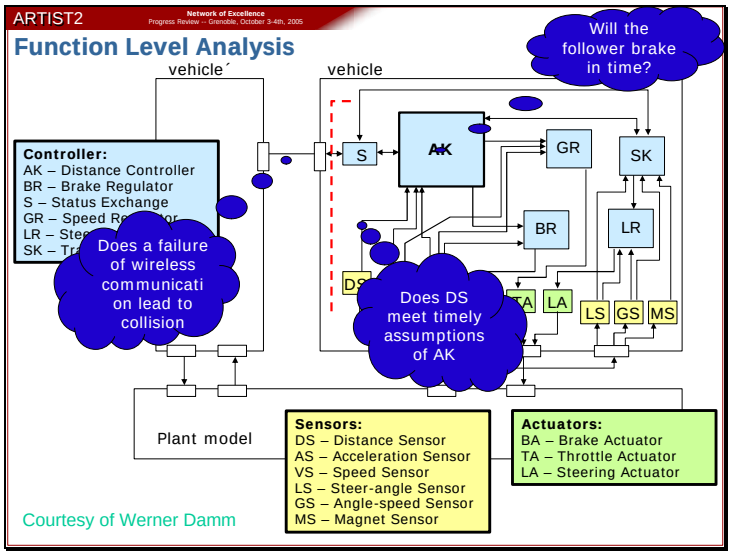

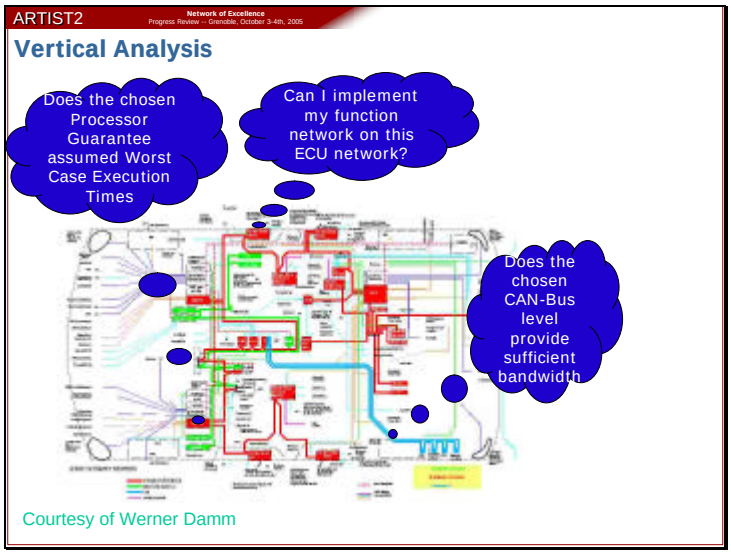

From these needs the following requirements emerge, for components of embedded systems:

- Components must be seen as part of *open systems*; they should support interface-based composition and refinement;
- Functional and non-functional aspects must be *jointly handled*, at both component- and (sub)system-level;
- Design space involves both functions and execution infrastructure;
- All this must be supported while dealing with *heterogeneous and flexible* Models of Computation and Communication (MoCC)

These requirements lead to the research challenges illustrated and explained in the following three figures:

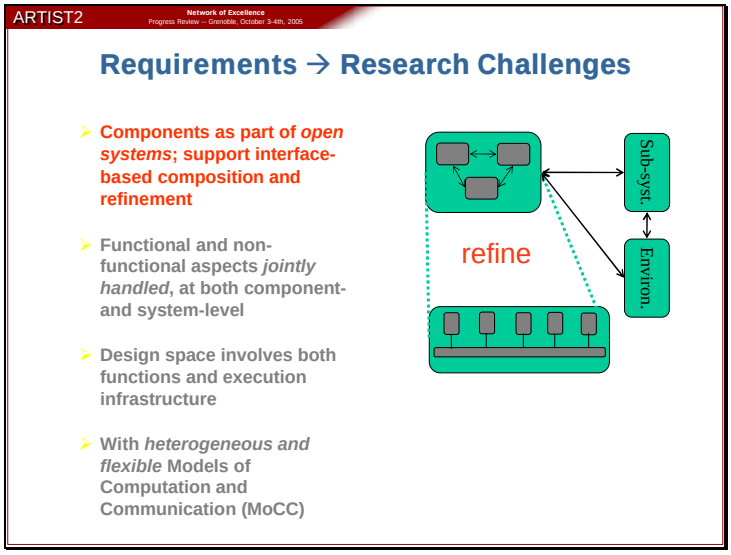

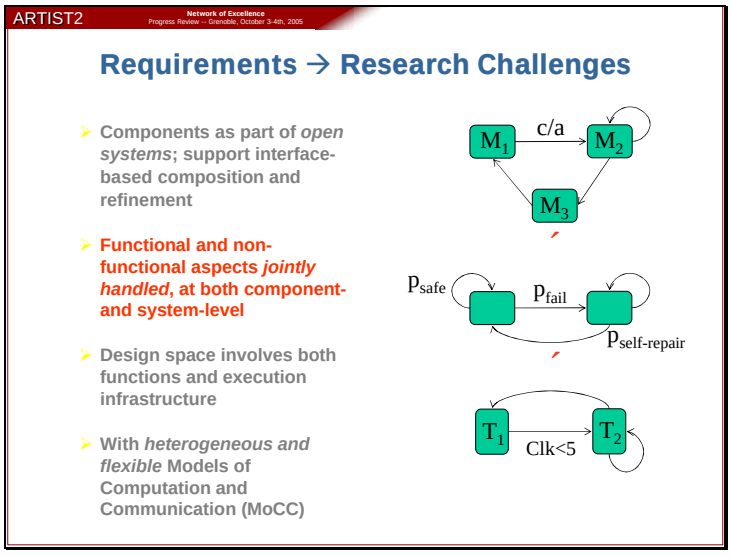

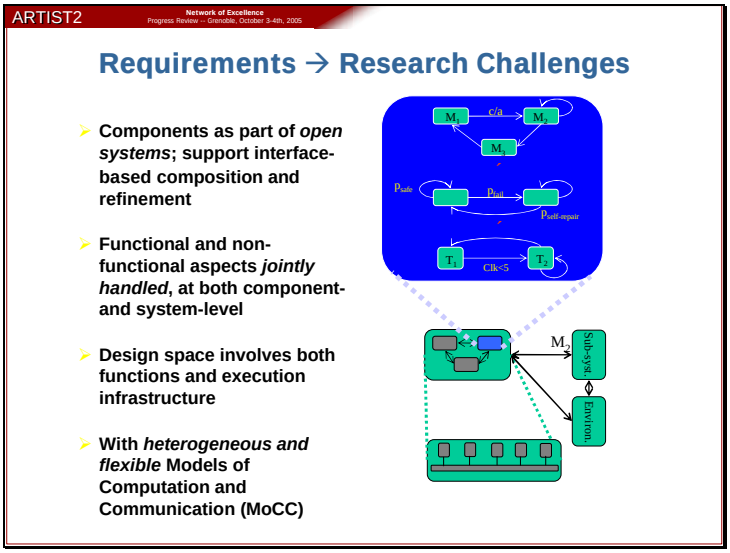

3.3 Milestones

We see the executive summaries of these two meetings as providing new avenues for research. And we regard these findings as the essential contribution of this ARTIST2 cluster.

3.4 Main Funding

Here we do not refer to the support needed to cover research activities that are ongoing at partners, but only to the support needed to perform the activities reported above. Corresponding main sources of funding are ARTIST2-HRT and other cluster's funds to support for participation to the various meetings and inviting affiliate partners.

Other funds than just travel used by partners correspond to the meeting preparation, contribution to the meeting minutes, and preparation of the material presented at the meetings – this material was often ad-hoc and not just standard reuse.

Most noticeably, industrial participants not members of ARTIST2 (and not affiliates) did not receive any support for their attendance to the Rome meeting. This was paid on their company's funds.

3.5 Indicators for Integration

We see the above results as a clear proof of team work. We think that:

- *The above results could not have been obtained by just standard interaction by attending conferences*. Face to face discussions in conferences and other usual meetings are typically much thinner in focus and less structured regarding dissemination effect.
- *The above results are different from the ones obtained in research projects, including other types of EU projects*. We do not see, e.g., STREPS or IPs spending such a large percentage of their effort in seeking for new research directions.

3.6 Evolution

The reader is referred to the next 18-month workplan for this point.

3.7 Interaction, Building Excellence Between Partners

Two meetings have been held gathering industrials and academics, one in Rome (January 2005) and Rennes (June 2005) which have allowed to discuss needs and solutions.

Existence of the group has lead to spreading new research directions in the embedded systems community, as can be seen from conferences and workshops in the area, e.g., Emsoft, ACSD, MemoCode.

Interactions within the group and across different clusters are summarized in the following two figures:

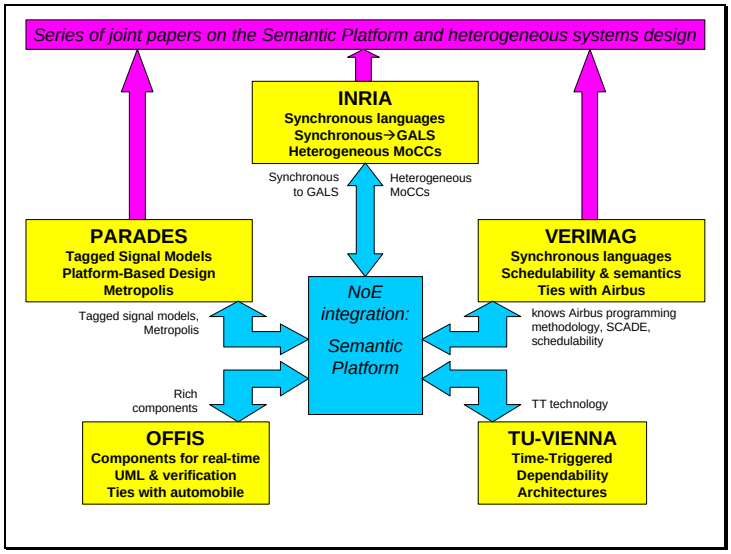

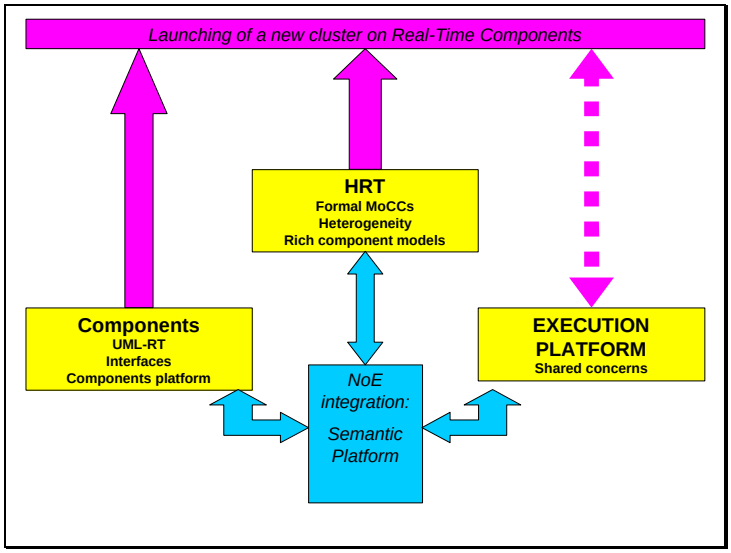

In the second figure, the dashed double arrow indicates tight interaction that contributed to the merge of the other two clusters – see the next 18 month workplan; the Execution Platform cluster, however, is not participating to the merge.

Here follows a list of joint papers that reflect ongoing cooperation:

INRIA-PARADES-VERIMAG: the following list of joint papers shows deep long-term cooperation. This cooperation was launched during the COLUMBUS project and was further sustained during ARTIST and ARTIST2:

- Benveniste, B. Caillaud, L. Carloni, P. Caspi, A. Sangiovanni-Vincentelli. ``Heterogeneous Reactive Systems Modeling: Capturing Causality and the Correctness of Loosely Time-Triggered Architectures (LTTA)''. *Proc. of EMSOFT'2004,* G. Buttazzo and S. Edwards, Eds., Sept. 27-29, 2004.
- Benveniste, L. Carloni, P. Caspi, A. Sangiovanni-Vincentelli. ``Heterogeneous Reactive Systems Modeling and Correct-by-Construction Deployment''. *Proc. of EMSOFT'2003,* R. Alur and I. Lee Eds., Oct. 2003.
- Benveniste, B. Caillaud, L. Carloni, P. Caspi, A. Sangiovanni-Vincentelli. Heterogeneous Reactive Systems Modeling: Capturing Causality and the Correctness of Loosely Time-Triggered Architectures (LTTA). In Proceedings of EMSOFT'04, LNCS, Sept. 2004.
- Benveniste, B. Caillaud, L. Carloni, A. Sangiovanni-Vincentelli. Tag Machines. In Proceedings of EMSOFT'05, LNCS, Sept. 2004.
- S. Prudhomme, W. Damm. Controlling Speculative Design Using Rich Components, Plenary *Address,* ACSD 2005. *(Note: S. Prudhomme is with Airbus France)*
- W. Damm, A. Votintseva, A. Metzner, B. Josko, Thomas Peikenkamp, Eckard Bode. Boosting Re-use of Embedded Automotive Applications through Rich Components. Foundations of Interface Technologies, FIT 2005 *(Note: this is not a joint-paper in the sense that authors are from different organisations. On the other hand, it is based on deep discussions with people from industry as written in the acknowledge section.)*

3.8 Spreading Excellence

The tight cooperation between academic and industrial partners allows solutions to spread quickly toward end-users via the industrial marketing services. This is the case of Verimag and Esterel, PARADES and ETAS, as well as TU Vienna and TTTech, or TU Munchen and BMW.

The following affiliate partners have been invited and supported by ARTIST2 for the two meetings:

- TU Munich: Jan Romberg
- Politecnico di Torino: Luciano Lavagno
- University of Udine: Tiziano Villa Villa@uniud.it
- University of California at Berkeley: A. Pinto (PhD Student)
- University of L'Aquila: S. Di Gennaro,

Note that the following industrial partners paid themselves for participating to the Rome meeting:

- GM: Tom Forest, Arnold Millsap
- BMW: Josef Berwanger, Tillmann Schumm

Finally, the affiliate partner

- U. Singapore: P.S. Thiagarajan

visited INRIA for 1 week on spring 2005, paid by ARTIST2 mobility funds.

4. Detailed Technical View

In this section, we gather general remarks on the topic. The below mentioned research cannot be acknowledged to ARTIST2 and is entirely funded by other means.

4.1 Brief State of the Art

European automotive and aeronautics industries are experiencing an exponential growth in functionality, with a drastic increase of innovations realized in software. The paradigm is shifting from an original “1 function = 1 ECU = 1 supplier” partitioning to distributed reali-zations of functions across multiple ECU’s involving multiple suppliers.

Several de-facto standards such as CAN, Flexray and the various OSEK extensions have found there way into series development. Model based development is increasingly gaining momentum, often involving automatic code generation. These processes are reaching substantial levels of maturity for single ECU implementations, including advanced capabilities for rapid prototyping and (hardware-in-the-loop) testing.

However, system-oriented design and virtual integration are supported only weakly, leading to late detection of integration problems. The OEM-supplier relation is mainly relying on textual initial requirements and well-established processes of delivering increasingly mature prototypes. Late requirement changes, incomplete initial requirements, or even inconsistent requirements are often leading to late design iterations or changes, with high incurred costs.

Already today, privately funded key initiatives like Autosar demonstrate the commitment of this industrial sector to reducing costs. Way to this is by harmonizing platforms and decoupling functional architecture design from target platforms.

As an interesting example of a major scientific difficulty, we recently learnt that the Autosar consortium will complete its work early 2006, without having converged on a *model for timing.* This is of course a central difficulty, since integrating subsystems from OEMs requires agreeing on how timing and schedulability are handled. This topic has been identified as a true research topic by the consortium.

On a similar line, A. Benveniste was invited to participate to the panel session at IEEE-Control and Decision Conference Dec. 2005: *How do control system design engineers use models and simulation?* Organized by Pieter J. Mosterman, from The Mathworks. The text of this panel session says (quoted, specific sentences underlined by us):

> In control system design, we typically model the plant in detail and then make the model amenable to control law synthesis. With this law at its core, the controller model is gradually refined with implementation detail. Physical models are combined with computational models to ensure we can realize the design. At present, computational modeling increasingly replaces physical modeling. This requires sophisticated modeling formalisms and tools. For example, in plant modeling, domain specific languages for, e.g., multi-body systems and image processing systems as well as extensive tool infrastructure, are needed. The challenges we face to further this trend are (i) providing domain-specific modeling formalisms, (ii) providing tool support, (iii) combining different formalisms, and (iv) automatic model translation. We discuss the role of models in control system design and address questions such as: Is there a set of sufficient semantic notions for our modeling languages or a general ‘computing API’ to combine different formalisms? Is simulation a sufficiently powerful technology? What is the best approach to generating modeling formalisms (libraries, meta-modeling, API, other)? Is there an optimal formalism to translate between formalisms? Can we derive

denotational or operational models from axiomatic specifications (i.e., generate models from 'scenarios')? How about producing target specific code? How can style guidelines be enforced and is there a need to configure tools for controller design? How about support for enterprise-wide modeling? Can model reduction techniques handle industrial models for control synthesis? How can you guarantee model composability? How can we obtain explicit models (e.g., hybrid automata) from models in a more practicable representation?

This text expresses very well concerns from industry. Note that these are specifically addressed by the HRT cluster and this JPRA.

On the other hand, state-of-the-art contributions from the academic circles are not numerous, and in fact, they come mostly from participants to this activity. The system-level design environment Ptolemy II supports component-based heterogeneous modeling and simulation of systems. It uses Henzinger-de Alfaro *interface automata* for the dynamic behavior specification. The Metropolis meta-model allows capturing non-functional aspects of design by so-called quantity managers, and provides means for declarative specification of non-functional constraints through its constraints logic. Recent work of Benveniste, Caillaud, Carloni, Caspi, and Sangiovanni-Vincentelli provides the foundations for heterogeneous reactive systems modeling and will help defining viewpoint synchronization and parallel composition of components involving heterogeneous viewpoints and Modes of computation and control.

4.2 Industrial Needs and Experience

Dealing with heterogeneous designs is essential to all industries developing systems with multi-skilled teams, methods, techniques, and tools. This includes:

- Avionics and aeronautics (Event and Time-triggered systems are developed and effectively used at Airbus Industries);
- Automotive (drive-by-wire, brake-by-wire);
- Rail Transport;
- Energy Production.

The following trends were identified by our industrial partners at the Rome meeting:

- GM technical strategy: reuse ECUs when possible (all vehicles, all model years), with no artificial bounds, no dependency on forecast of future, implies immunity to deployment platform differences. Separate logical feature content from physical deployment platform (*this refers to non-functional characteristics*); use standard control infrastructure SW, and deploy OTS feature content to optimal physical architectures.
- Control representation is provided. Functional partitioning is a key issue. Deploy functional features on physical ECUs; this partitioning may differ from vehicle to vehicle. Perform both logical and ECU reuse as much as possible. GM wants to allow both reuse of features and ECUs. Of course, there are communication differences due to partitioning. Need to cope with this difficulty without using adaptive middleware.

4.3 Ongoing Work in the Partner Institutions

- Albert Benveniste and Benoît Caillaud (INRIA), Paul Caspi (Verimag), Luca Carloni (Columbia Univ. in the USA), and Alberto Sangiovanni-Vincentelli (PARADES) work jointly since now 4 years on developing a comprehensive theory of heterogeneous systems and their heterogeneous composition. This theory builts on the original tagged

systems model proposed in 1998 by Lee and Sangiovanni-Vincentelli in support of the Ptolemy tool. Very much like all composition operations on systems can be subsumed by more abstract concepts defined on category theory, this group of people have defined an abstract model of system in which the Model of Computation and Communication (MoCC) can be seen as a parameter and can be adjusted at needs. The ultimate objective is to lift up deployment analysis to a generic problem, where algorithms and tools can be developed at an abstract level and instantiated for the different MoCCs encountered and their combination. This is long way effort and one can say that about 60% of the research has already been done. It is expected that this will result in techniques for deployment analysis that will better scale up than just relying on existing low level expansion of deployment models and their model checking.

- Suzanne Graf, Joseph Sifakis (Verimag) and Gregor Goessler (INRIA) have developed a model of real-time component where some properties such as deadlock freedom are preserved by construction. The uniform scheduling technique used is by priorities. This model has been more recently lifted to encompass flexible MoCCs, not only timing.
- The Metropolis effort at University of California, Berkeley and PARADES constitutes a framework where models of different MoCCs can be combined. Their different MoCCs are described by means of a uniform low-level model of interface automata. Non-functional characteristics are captured by means of *quantities*. Metropolis supports design space exploration and platform-based design.
- The group at OFFIS has developed a new theory of rich component models, which is rich enough to support both functional and non-functional characteristics. This theory addresses requirements on component models supporting library based development processes of embedded automotive and avionics applications. A key requirement for such a component model is the need to cater for multitude of non-functional constraints (including e.g. resource constraints, real-time requirements, and safety requirements). These research activities benefits from deep discussion with key persons from the automotive and avionics industries.

4.4 Main Funding (not ARTIST2)

Main sources of funding include (the list below is not comprehensive):

- Swiss National Science Foundation (Analytic Estimation Methods)
- DECOS IP
- IST project COLUMBUS
- IST project RISE
- AUTOMODE Joint project