ARTIST 2
Network of Excellence

IST-004527 ARTIST2:
Embedded Systems Design

Activity Progress Report for Year 1

JPRA-NoE:

Merging the Event-Triggered and Time-Triggered Paradigms

Cluster:

Hard Real-Time

Activity Leader:

Paul Caspi (Verimag)

Two approaches for designing and implementing synchronous hard real-time systems have been developed: event triggered and time-triggered. These have been applied separately to large-scale embedded systems in Europe.

This activity will investigate how to integrate these approaches in a semantically sound and efficient design flow, preferably platform-based. It will require in particular research on mathematical modelling and simulation of both time-triggered and event-triggered architectures, research on RTOS execution mechanisms and research on performance evaluation and optimisation.

This is fundamental work on merging two of the main paradigms in real-time systems design. The expected results are important from both a theoretical point of view, and also for industry (distributed embedded systems and network on chip applications).

Table of Contents

1. Introduction

1.1 Activity Leader

Team Leader: Paul Caspi (Verimag)
Areas of their team's expertise: synchronous languages, strong collaboration with Esterel Technologies and Airbus Industries.

1.2 Clusters

Hard Real-time
Adaptive Real-time
Execution Platforms

1.3 Policy Objective

Fundamental work on merging two of the main paradigms in real-time systems design. The expected results are important from both a theoretical point of view, and also for industry (distributed embedded systems and network on chip applications).

Two approaches for designing and implementing synchronous hard real-time systems have been developed: event triggered and time-triggered. These have been applied separately to large-scale embedded systems in Europe.

This activity will investigate how to integrate these approaches in a semantically sound and efficient design flow, preferably platform-based. It will require in particular research on mathematical modelling and simulation of both time-triggered and event-triggered architectures, research on RTOS execution mechanisms and research on performance evaluation and optimisation.

1.4 Industrial Sectors

Unifying the two paradigms is essential for systems mixing low-level continuous control with higher-level supervision, used in:

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

and, in general, most safety-critical control systems which require that these two paradigms be mixed in a semantically sound and manageable way.

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 and hardware industries, expertise in design flow for Hard Real-Time.

Team Leader: Albert Benveniste (INRIA)
Areas of his team's expertise: synchronous languages.

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

Team Leader: Petru Eles (Linköping University)
Areas of his team's expertise: schedulability analysis for heterogeneous systems.

Team Leader: Tom Henzinger (EPFL)
Areas of his team's expertise: development of abstract programming models for real-time computing [Giotto: time-triggered; xGiotto: both time- and event-triggered].

Team Leader: Rolf Ernst (University Braunschweig)
Areas of his team's expertise: formal performance models for networks-on-chip.

2.2 Affiliated partners and Roles

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

Team Leader: Thomas Thurner / Hermann von Hasseln (DaimlerChrysler)
Areas of his team's expertise: automotive industrial case study.

Team Leader: Stephan 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: Jan Romberg (TU München)
Areas of this team's expertise: Synchronous languages, model-based development, automotive applications.

Team Leader: Christofer Kirsch (University of Salzburg)
Areas of this team's expertise: development of abstract programming models for real-time computing [Giotto: time-triggered; xGiotto: both time- and event-triggered].

2.3 Starting date, and expected ending date

Starting date: September 1st, 2004

Expected ending date: September 1st, 2006

2.4 Baseline

In the past, real-time system design was dominated by the choice between two paradigms: event triggered (ET) and time triggered (TT). In ET systems, activities are initiated by events, such as interrupts, whereas in TT systems, activities are triggered by time-dependent events.

ET and TT have been considered to be distinct and incompatible paradigms that are difficult to combine within a single distributed architecture. While the TT approach allows guaranteeing dependable temporal performance, it often makes inefficient use of resources and is not sufficiently flexible in handling dynamic service requests.

The ET approach provides more flexibility. Many current applications follow neither the TT nor the ET scheme entirely, combining features of each.

We plan to provide methods for the efficient combination and cooperation the TT and ET approaches.

2.5 Technical Description

Currently, the combination of the ET and TT paradigms has only been handled in the specific situation where the ET tasks are less urgent than the TT ones. In this case, an ET tasks is allocated to idle TT slots. Solutions where ET and TT activities have equal priorities are more difficult to implement. These call for:

- Scheduling and execution mechanisms that are coherent with the above (we want to execute what has been simulated and validated).
- Formalisms and methods for simulation and verification of the mixed ET and TT architectures.
- Methods for simulating, evaluating and optimising performance.

These subjects will be developed in tight cooperation between the participating clusters

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 Semantic Platform.

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 noticed by car 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; 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. GM prefers to reuse ECU that are statically configured once for all possible vehicles rather than dynamically adapted to the target vehicle. To accomplish this GM tries to use communication strategies that make the ECUs as immune to vehicle-to-vehicle communication differences as possible.

If CAN messages are used, even if the two ECUs receive the same data, the data are received packaged differently (in different frames). This induces platform specific topology that is a barrier to reuse.

Therefore, chosing between or combining smoothly the different ET and TT paradigms is a key issue in core architecture design.

From these considerations, the following research directions have been suggested by our GM participant; these have been validated by the BMW participants too and then extensively discussed in a 2 hour forum by all participants:

- Desire to use ECU reuse → preference for techniques that make applications *insensitive* to changes rather than provide automatic adaptation techniques.
- Develop techniques to ensure incremental addition of functions so that architecture is changed minimally.
- How to select an architecture that is optimal in scalability and extensibility?
- Scheduling:
    - Develop techniques to decouple application schedules from communication schedules, in a reasonably optimal way.
    - Select *best schedule* from certain metrics among possible ones. Even determining *what the metrics should be.* Find metrics that reflect extensibility and scalability.
    - Is it possible to have static scheduling generation for distributed implementations? One of the problems is with the suppliers: you can calibrate a schedule a priori, but when the supplier changes something, then this schedule is no longer valid.
- How can we apply the above techniques with *incomplete and approximate information*, for early architecture decisions. (Much harder to do with TT than it was with ET CAN in the past.)
- Do we have tool support for the above issues? This is particularly needed for distributed application → distributed deployment & scheduling.

3.1.2 Other important findings and notices

- Paul Pop (Linköping, Executions Platform cluster). Partitioning between ET and TT modes for a given set of tasks is an important issue. At the moment, the standard way provided by, e.g., Flexray bus, is by selecting a bandwith limit such that TT is taken for the higher band whereas ET is taken for the lower one. This may not be flexible enough.
- Paul Caspi (Verimag): deploying tasks over ET/TT architectures may raise issues of preservation of semantics. Usually, engineers pay a great deal of attention to schedulability issues and their incremental nature, but are not aware of this issue of semantics preserving. Therefore this can cause a mismatch beween the functions being specified and the actual implementation. This is considered a novel and important direction for research.

- Roman Obermaisser (TU Vienna) and Yves Sorel (INRIA): one interesting way of achieving good integration of ET and TT is by emulating one scheme on top of the other. For example, one may try to emulate ET on top of TT (TU Vienna), or the converse (INRIA). These proposals raised a number of hot discussions regarding how far such emulations can be considered correct and comprehensive.

3.1.3 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.

Still, the findings of the Rome meeting could be refined in the Rennes meeting by the presentations of Roman Obermaisser (TU Vienna) and Lothar Thiele (ETH, Execution Platform cluster). This is summarized next.

3.2.1 Updating the findings from Rome meeting

- *Summary of discussion that was held with GM and BMW engineers at the Rome meeting regarding emulating ET on top of TT, and update:*
    - There was a big debate at this Rome meeting. This is why Roman Obermaisser insisted in his talk on the exact/approximate emulation of CAN on top of a time-triggered architecture. The focus of the discussion at the Rome meeting was whether an event-triggered communication service on top of a time-triggered architecture “is still CAN” when there is no CSMA/CA media access control strategy at the physical layer. The answer is that the integrated architecture for CAN-based and TT applications supports (optionally) the execution of the CSMA/CA media access control strategy on top of the TDMA scheme of the core architecture. This execution of CSMA/CA is called protocol emulation and only necessary for legacy applications. For legacy applications, protocol emulation ensures that existing CAN-based legacy software works correctly in a time-triggered architecture. With protocol emulation a virtual CAN network exhibits the same temporal message order as a physical CAN network.
    - Yes, there is an overhead resulting from this emulation if full emulation is wanted.
- *A more systematic approach to schedulability that is suitable to component based design.* Such an approach was presented by Lothar Thiele, based on so-called real-time calculus, a technique to deal with dates of events in an algebraic and compositional way. How to widen the scope of this approach (today restricted to control-independent schedules) is an intesresting direction for research.

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 thiner 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.

The interaction between partners in this group is best described by using the following two figures, which indicate the interaction within the cluster and with other clusters.

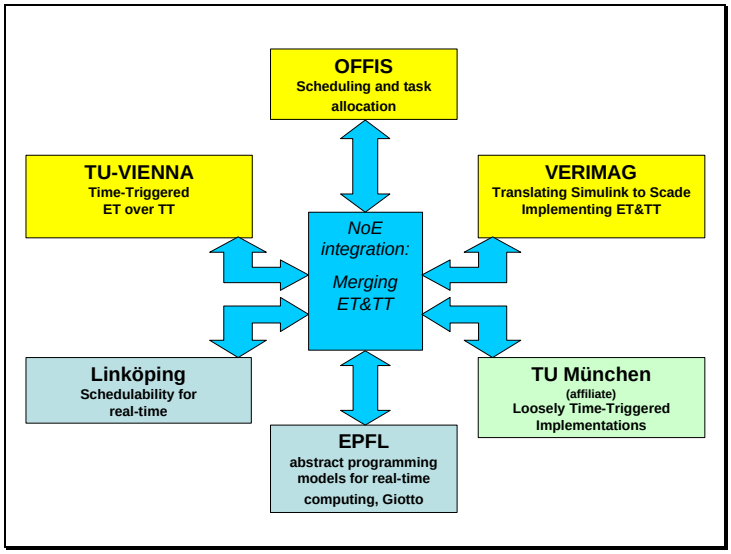

The boxes in grey indicate partners belonging to other clusters (components and execution platforms) and the green box indicates an affiliate partner from outside ARTIST2.

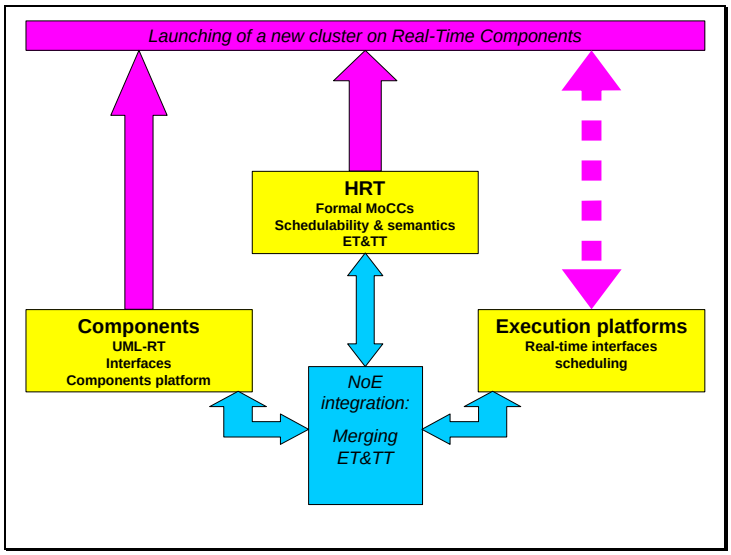

For the moment, only informal cooperation between partners has been triggered:

- Parades was influenced by Verimag in his work on ET&TT implementation. It resulted in two separate papers on the same subject being presented at Emsoft05. A still informal but tighter cooperation has been decided to be untertaken since then.
- Some provisions toward launching a common Strep proposition were taken between Verimag and TU München. However it failed because of the competing positions of their respective “natural” indutrial partnersEsterel Technologies and ETAS.

3.8 Spreading Excellence

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

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

Two approaches for designing and implementing synchronous hard real-time systems have been developed: event triggered and time triggered. These have been applied separately to large-scale embedded systems in Europe.

In ET systems, activities are initiated by events, such as interrupts, whereas in TT systems, activities are triggered by time-dependent events.

ET and TT have been considered to be distinct and incompatible paradigms that are difficult to combine within a single distributed architecture. While the TT approach allows guaranteeing dependable temporal performance, it often makes inefficient use of resources and is not sufficiently flexible in handling dynamic service requests.

The ET approach provides more flexibility. Many current applications follow neither the TT nor the ET scheme entirely, combining features of each.

Currently, the combination of the ET and TT paradigms has only been handled in the specific situation where the ET tasks are less urgent than the TT ones. In this case, an ET tasks is allocated to idle TT slots. Solutions where ET and TT activities have equal priorities are more difficult to implement. These call for:

- Scheduling and execution mechanisms that are coherent with the above (we want to execute what has been simulated and validated).
- Formalisms and methods for simulation and verification of the mixed ET and TT architectures.
- Methods for simulating, evaluating and optimising performance. These subjects will be developed in tight cooperation between the participating clusters.

Reseach is ongoing at the different partners to provide methods for the efficient combination and cooperation the TT and ET approaches.

This activity aims at investigating how to integrate these approaches in a semantically sound and efficient design flow, preferably platform based. It requires in particular research on mathematical modelling and simulation of both time-triggered and event-triggered architectures, research on RTOS execution mechanisms and research on performance evaluation and optimisation.

4.2 Industrial Needs and Experience

Unifying the two paradigms is essential for systems mixing low-level continuous control with higher-level supervision, used in:

- 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 was noticed at the Rome meeting:

- General Motors is considering the use of Flexray. There are different categories of communications: fault tolerant, closed loop control, backbone/bridge, body, infotainment, field bus, legislated diagnosis… FlexRay is best suited for the first 3 categories. GM feels that the flexibility in FlexRay does not come from the dynamic segment; feel this segment is rather inefficient. It's difficult to make use of the end of dynamic segment portions remain unused. Latency analysis is difficult. The effects on the dynamic segments on asymmetric faults are significant and persist for the remainder of the dynamic segment. Some aspects of FlexRay that do support reuse. Devices can send in more than one slot in the static segment. Allows the communication in a static slot to be treated as the communication of a Virtual Device and VD can be moved. TDMA offers the property that communications in slots other than those sent or received by a node do not affect the node. Another aspect is the use ofmessage constructs to indicate protocol relevant characteristics. GM needs techniques that decouple application schedules from communication schedules. One way of doing this is not to use the fine grained structure of timing. Communications sent sometime during a cycle – data received in the last cycle isused by the application in the current cycle. But this has drawbacks as well since data dependencies introduce multi-cycle delays. Decoupling techniques can be a very valuable research problem. TT comm. protocols are capable of microsecond level synchronization of tasks. Is this needed? Not obvious. In many cases it seems that the protocol itself drives the requirement of high accuracy synchronization, not the application. The inherent Composability of TT is desirable. From a flexibility perspective, very fast ET protocols would help better (no scheduling issues). Problem is that physics may prevent a fast prioritized ET protocol- high speed may require TT. GM doesn't want to underestimate the move to TT. This requires the application to be synchronized on the communications, not the other way around – very different from the way it is done today. TT protocols behave very differently than CAN; retransmissions take more time. The strategy of how to use TT needs to be learnt and understood. It is difficult to coordinate TT schedules on a system-wide basis. Much more difficult than in ET systems.Significant effort is needed to develop infrastructure necessary for a new protocol
- Josef Berwanger at BMW works on introducing FlexRay at BMW: Benefits of FlexRay technology: high bandwidth (x 25), determinism in static segment, task synchronisation in distributed systems, short cycle times, reliable communication, enabling system integration, extensible, makes it possible to implement X-by-wire functions. Synchronisation of tasks for sensors, control functions, and actuators becomes possible. Cycles can be exactly repeated with their synchronisation; strict periodicity can be guaranteed. FlexRay allows for Composability in the time domain. System integration is easier if components use different slots of the bus integration by simple interleaving. Allowed slots for each type of component is maintained in a central data base to ensure this. Data can be repeated within the same cycle (tc0). This gives the possibility to define different cycles for the application. The dynamic part is used for ET signals and transport layer (e.g., for diagnosis or flash data download) where it is definitely more efficient.

4.3 Ongoing Work in the Partner Institutions

- Jan Romberg (TU Munich) in the AUTOMODE Joint Project with Validas AG, Bosch, BMW, ETAS works on loosely time-triggered implementation scheme for CAN : Network nodes are logically arranged in a tree. Synchronizing links have to have periodic traffic with base period T (put don't cares if needed). Synchronous dataflow programs are mapped to the cascade. A synchronizing message triggers node activation; this propagates throughout the cascade. Send/receive phases alternate with computation phases, at each node. Non-synchronizing messages are buffered. Timeouts are set and tuned according to jitter estimates. Self-triggering by timeouts can trigger the reading of messages from buffer and launch the computation. If losses are temporary, resynchronization can take place and yields robustness against intermittent losses.

- Roman Obermaisser (TU Vienna) works within the Decos project on TT-ET Integration: CAN on top of TTP/C: CAN is an example of ET protocol, widely used and low cost. Disadvantages areknown: variability in latencies, limited throughput, no atomic multicast, no handling ofbabbling idiot failures. Physical integration, good for mixed-criticality applications, allows legacy reuse, and improves the quality of CAN communication services.TT platform is used to put both TT and CAN based services on top of it. The CAN that is put on top does not invalidate the certified properties of TT layer. High-level CAN services are implemented by dividing TT segments into a TT slot followed by a CAN segment (looks like FlexRay division). The interest is that this gives flexibility for how much you allocate to CAN: tunable. The CAN arbitration protocol isemulated by "modeling" how CAN would behave.ET protocol performed by emulation. Temporal performance demonstrates that it is indeed interesting not to emulate exactly CAN ordering of messages: one can do better in terms of latency and jitter. Disabling this reordering can be acceptable for some applications; this yields a "modern" ET solution. Another interest is the improvement with respect to fault tolerance.

- Paul Pop from Linkoping works with Volvo on the scheduling and optimization of time- and event-triggered distributed embedded systems. The problem is to perform partitioning between ET and TT scheduling for given tasks. This has an impact on schedulability since preemption is permitted in one case but not in the other.

- Alexander Metzner (OFFIS) works on modeling combined event and time triggered systems for automatic allocation of distributed task systems. Incremental integration requires task and message budgeting. Jitter is a difficulty for incremental integration; safe overapproximation is used for that. Allocation is determined using SAT checking techniques from the formal verification area. SAT checking checks for satisfiability of mixed integer/Boolean formulas. If this procedure terminates, "nearly-optimal" allocation is achieved (because of overapproximation). Computation times for this are highly variable. This schedulability analysis techniques extends for mixed RTOS as well, e.g., OSEKtime.

- Paul Caspi and Stavros Tripakis (Verimag) work with Esterel Technologies in integrating model-based design and preemptive scheduling in mixed TT and ET systems within the SCADE environment. Deterministic semantic preserving communication schemes relating ET and TT tasks have been obtained and are implemented in SCADE.

- A.Sangiovanni-Vincentelli, A. Ferrari, L.Mangeruca and M.Baleani (PARADES GEIE) are working on a model-based software design flow aiming at automatically mapping synchronous models, as provided by tools such as Ascet, SCADE, ESTEREL and Simulink/Stateflow (the latter with some restrictions). In particular, the implementation of buffer techniques and definition of scheduling constraints (e.g. synthesis of scheduling priorities) are investigated.

4.4 Main Funding (not ARTIST2)

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

- Swedish Foundation for Strategic Research (SSF)
- Centre for Industrial Information Technology (Linköping University)
- Swiss National Science Foundation (Analytic Estimation Methods)
- DECOS IP
- IST project COLUMBUS
- IST project RISE
- AUTOMODE Joint project