



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

Activity Progress Report for Year 3

Scheduling and Resource Management

Cluster:

Operating Systems and Networks

Activity Leader:

Professor Alan Burns (University of York, UK)

<http://www-users.cs.york.ac.uk/~burns/>

Policy Objective (abstract)

The management and scheduling of system resources is one of the main development challenges in any embedded systems. This activity is concerned with multi-resource policies and analysis techniques that allow safe but effective resource utilisation. All resource types are considered: processing units, communication units, storage units, application-specific units and generic resources such as power.

Versions

number	comment	date
1.0	First version delivered to the reviewers	January 15 th 2011

Table of Contents

1.	Overview of the Activity	3
1.1	ArtistDesign Participants and Roles	3
1.2	Changes w.r.t. Y2 deliverable.....	3
1.3	Affiliated participants and their role within the Activity.....	4
1.4	Changes w.r.t. Y2 deliverable.....	4
1.5	Starting Date, and Expected Ending Date	4
1.6	Policy Objective.....	5
1.7	Background.....	5
1.8	Technical Description: Joint Research.....	6
1.9	Work achieved in Year 1 (Jan-Dec 2008)	7
1.10	Work achieved in Year 2 (Jan-Dec 2009)	8
1.11	Problems Tackled in Year 3 (Jan-Dec 2010)	10
2.	Summary of Activity Progress in Year 3 (<i>Jan-Dec 2010</i>)	12
2.1	Technical Achievements	12
2.2	Individual Publications Resulting from these Achievements	19
2.3	Interaction and Building Excellence between Partners	25
2.4	Joint Publications Resulting from these Achievements.....	26
2.5	Keynotes, Workshops, Tutorials.....	28
3.	Milestones, and Future Evolution	29
3.1	Problems to be Tackled in Year 4 (Jan 2011 – Dec 2011).....	29
3.2	Current and Future Milestones	29
3.3	Main Funding	30
4.	Internal Reviewers for this Deliverable	31

1. Overview of the Activity

1.1 ArtistDesign Participants and Roles

Professor Alan Burns (University of York - UK)

The Real-Time Systems Research Group at the University of York contributes research on advanced scheduling and resource management policies.

Professor Giorgio Buttazzo (Scuola Superiore Sant'Anna -Italy)

The Scuola Superiore Sant'Anna (SSSA) of Pisa investigates advanced scheduling methodologies for increasing the predictability of real-time systems characterized by a highly variable workload and execution requirements.

Professor Luis Almeida (University of Porto - Portugal)

The team at the University of Porto (UnivPorto) is involved in the design and analysis of tools and mechanisms for supporting dynamic QoS management, mainly for distributed multimedia systems, flexible scheduling, dynamic reconfiguration, graceful degradation and survivability for distributed embedded control systems, particularly robots and vehicles.

Professor Michael Gonzalez Harbour (University of Cantabria – Spain)

University of Cantabria focuses on the integration of the resource management techniques developed by the other partners in the integrated framework for flexible resource management (FRESCOR). The group also participates in the development of the Real-time POSIX operating systems standards and the OMG standard for Modelling and Analysis of Real-Time Embedded Systems (MARTE).

Professor Gerhard Fohler (University of Kaiserslautern - Germany)

The Technical University of Kaiserslautern (TUKL) works on the integration of offline and online scheduling for combining time triggered and event triggered methodologies in the same system and provide resource management methods for media processing.

Professor Karl-Erik Årzén (University of Lund - Sweden)

The team at Lund University (ULUND) works on scheduling of embedded controllers, in particular co-scheduling approaches, and the use of feedback approaches in resource scheduling.

Professor Eduardo Tovar (Polytechnic Institute of Porto – Portugal)

The team at the Polytechnical Institute of Porto is involved in Scheduling on Multicores, QoS-Aware in Distributed and Collaborative Computing, Resource Management in Sensor Networks and general purpose abstract models and dynamic run-time adaptability with anytime approaches.

Professor Maja D'Hondt (IMEC)

The runtime resource management team at the Digital Components division in IMEC is focusing on task scheduling, data storage and access methodologies to improve the performance and energy consumption of dynamic software applications, running on MPSoC platforms.

1.2 Changes w.r.t. Y2 deliverable

Stylianos Mamagkakis has been replaced by Maja D'Hondt in IMEC during the current year.

1.3 Affiliated participants and their role within the Activity

Professor Alfons Crespo (Technical University of Valencia – Spain, Affiliated to Cantabria)

The team at the Technical University of Valencia is involved in providing real-time memory management OS support, and real-time kernel virtualization.

Professor Marisol García Valls (Carlos III University of Madrid - Spain, Affiliated to Cantabria).

The team at the Carlos III University of Madrid works on building real-time support into middleware for embedded systems and memory-based QoS management techniques to provide support for predictability in Real-Time Java middleware.

Professor Alejandro Alonso (Technical University of Madrid – Spain, Affiliated to Cantabria)

The team at the Technical University of Madrid investigates on integrated resource management policies with emphasis on adaptability.

Professor Lucia Lo Bello (Technical University of Catania - Affiliated to Pisa)

The team at the University of Catania works on QoS-oriented scheduling and management of communication and processing elements in embedded platforms, including energy-aware solutions.

Professor Pau Martí (Technical University of Catalonia – Affiliated to Lund)

The team at the Technical University of Catalonia works on the integration of feedback control and resource management techniques to provide adaptability to changing conditions on both resource and applications demands.

Professor Tullio Facchinetti (University of Pavia - Affiliated to the Scuola Superiore Sant'Anna)

The University of Pavia considers new methodologies for integrating overload management techniques with energy-aware strategies, in the context of small embedded systems for battery operated devices.

Professor Hermann Härtig (University of Dresden – Affiliated to Kaiserslautern)

The team at the University of Dresden are involved in building micro-kernel- and hypervisor-based systems as experimentation platforms.

Professor Paulo Pedreiras (University of Aveiro – Affiliated to UnivPorto)

The team at University of Aveiro, previously a core participant, focuses on the development of scheduling and QoS management policies to cope with dynamic networked embedded systems.

1.4 Changes w.r.t. Y2 deliverable

Dr Steffens from NXP has retired, with no replacement as yet.

1.5 Starting Date, and Expected Ending Date

This activity started at the commencement of ArtistDesign and continues work undertaken in the previous ARTIST NoE. Although a number of milestones are expected to be achieved and reported on during the duration of this NoE, scheduling and resource management will always be a research focus as the nature of the resources change and the needs of applications expand. The research topic will therefore extend beyond the lifetime of ArtistDesign.

No changes with respect to Year 2.

1.6 Policy Objective

The main objective of this activity is the provision of models of embedded platform resources and policies, and the necessary analysis for undertaking the run-time scheduling of these resources and policies. A key scientific challenge is to link this resource-centred analysis with models of the application (and their resource usage policies) and the performance profiles of the hardware platform itself. Issues of temporality, safety, reliability and security can only be effectively addressed by an integration of these various abstract views of the overall system.

Seven promising approaches for providing this integration are:

- the use of search techniques to investigate architectural tradeoffs,
- the definition and use of virtual (unshared) resources,
- the use of reservations and contracts to allocate virtual resources,
- the use of coordination languages to integrate the use of different resource types,
- taking advantage of parallel processing platforms, such as multicores and FPGAs, in order to satisfy timing requirements,
- the application of self-adapting (feedback) resource allocation algorithms, and
- the recognition of the various time scales over which resource management must occur.

The nature of the scientific challenge should not be underestimated. Although very effective results for single resource (e.g. the processor) scheduling are available (and are used in industrial practice), for multiple resources there are no current applicable theories that have wide acceptability. Even for multi-processor SMP systems there is no consensus on the appropriate means of managing this resource.

The impact on operating systems will be taken into account via interactions with Activity 1 of this cluster. In addition the management of the network resource(s) will be addressed via joint work with Activity 3.

The industrial domains that will directly benefit from the results of this research include consumer electronics (in particular the games industry and multimedia applications), the automotive and aerospace industries, and environmental electronics such as smart spaces.

No changes with respect to Year 2.

1.7 Background

The platforms on which the next generation of embedded systems will be implemented will be radically different from those used in the current generation. The scale, performance, scope and applicability are all subject to significant enhancement. This presents the application developer and systems engineer with a number of fundamental challenges. At the centre of these challenges is the (effective) management of the platform's resources. Such platforms are likely to be multi-core (64 soon and 200+ within a few years); involve buses and networks of various capabilities and speeds (both off-chip and on-chip, i.e. NoCs); memories of various

speeds; include specialised components such as MEMS, ASICs, DSPs, and ASIPs; are linked to a wide variety of sensors and actuators; are embedded in systems powered by batteries (for mobile applications); include areas of FPGA (which are capable of dynamic reprogramming); and may have input/output links to global web-based information systems (for cyber-physical systems). Applications will be multi-resource and configurable. They will want to make dynamic modifications to their behaviour to support adaptability and environmental change. For example, the level of parallelism may alter at run-time and lead to re-evaluation of how this parallelism is delivered, e.g. by a subset of the cores, by application specific processing elements or by reprogramming an area of FPGA.

The main objective of this activity is to investigate how this wide variety of platform resources can be abstracted, modelled and managed, and application-specific resource allocation policies defined. At run-time, near optimal resource usage is desirable, but so are levels of protection for high integrity applications and those that have security constraints. Effective run-time scheduling of multi-resource platforms is not currently achievable; new methods will need to be developed.

No changes with respect to Year 2.

1.8 Technical Description: Joint Research

The technical achievements expected range from specific scheduling algorithms that cater for particular groups of resources, to a general purpose framework for addressing the broad problem of managing multiple resources for multiple applications on multiple time scales with multiple policies. It is expected that a means of abstracting, via a parameterised definition, the capability of each resource will be developed. A greater understanding of the distinctive roles of both static architectural tradeoffs and dynamic run-time adaptability will be obtained by both theoretical study and where possible the analysis of industrially relevant case studies.

The activity will focus on the techniques needed elsewhere in the NoE for predictability and adaptability. It will directly address the run-time techniques and analysis that will need to be supported by the OS and any network protocols.

The first 18 months will focus on producing an outline taxonomy on scheduling system resources and the analysis techniques available to manage their use. One aspect of this taxonomy will be to survey the various forms of parallelism becoming available on current platforms; other topics will be the use of hierarchical scheduling, “anytime” approaches and specialised hardware. For mobile platforms, energy is a key resource that is the subject of much research that will be surveyed. The final class of resources to be considered is that containing specialized components and external devices (and information sources).

It is expected that within 4 years, real-time scheduling algorithms for multicores with a utilization bound greater than 50% will be developed for sporadically arriving tasks. These results will be extended for arbitrary deadlines and for dealing with shared data structures.

Reconfigurability is a key issue for some applications. It is essential to not only ensure that the new mode is safe but also to ensure that the transition to the new mode does not violate timing requirements; this is often referred to as the mode change problem, and it is currently unsolved for multicores. Considering the current state-of-art in real-time scheduling in multicores, we expect this result on multicores to be available through the progress of ArtistDesign.

Dynamic memory management has been systematically avoided in real-time systems. One of the main reasons for this is the absence of deterministic allocators. Recently a new algorithm for dynamic memory allocation (TLSF) that solves this problem of the worst case bound whilst

maintaining the efficiency of the allocation and de-allocation operations has become available. This allows the reasonable use of dynamic memory management in real-time applications and permits consideration of dynamic memory as a first-class resource which can be used jointly with other resources in the schedulability of embedded systems. This integration of memory management and other resources is likely to become increasingly significant in the coming period.

No changes with respect to Year 2.

1.9 Work achieved in Year 1 (Jan-Dec 2008)

As indicated in the Description of Work the first milestone set for this activity during 2009 was to produce, in the second year, the outline of a **taxonomy** of resource management and usage. This taxonomy is aimed at being as broad as possible in the sense that:

- It covers all forms of processing devices from single processors to multicores, FPGAs etc.
- It covers all forms of communications including NoCs – although networking itself is covered by another activity within this cluster.
- It covers all specialised components; ASIC etc – although platform issues are covered more in another cluster.
- It covers system-wide resources such as energy.
- It covers all forms of (offline) verification including simulations, analysis (scheduling analysis), and various forms of model checking – although this later topic is covered more in another cluster.
- It covers online resource management (including control issues) to achieve openness, adaptability and fault tolerance.

Ongoing research on scheduling focuses on resource specific analysis, various resources have been considered including single processor with fixed priority or EDF scheduling, multiprocessor platforms of various types (including virtualised execution environments, multicore systems, and caching issues), communications media (but see Network activity report), memory and energy. Work has been done to extend the Contract Model (developed in the FRESCOR project) to include a wider range of resources and multiprocessor systems.

A focus on control applications has addressed period selection (sampling periods), low cost implementations and event-based control. Improvements to control effectiveness have been addressed by applying sensitivity analysis. Also the use of predictable CAN for control applications has been studied.

Another key problem addressed is the Adaptive Management of Multiple Resources. The resources typically used in end-to-end delivery of data streams often exhibit fluctuating availability and interdependencies. Wireless networks, for example, are influenced by interference, mobility, or physical structures, which cannot be known before system deployment. Even on single devices, a number of resources will be interdependent making the issue of multi resource management important. The focus of efforts has been put towards integrating CPU scheduling and cache management for efficient cache use and predictability.

The first year deliverable described technical achievement under the following headings:

- **Towards a taxonomy of resource usage**

- **Flexible scheduling framework**
- **Multi-resource scheduling on multicore platforms**
- **Memory arbitration on heterogeneous multicore platforms**
- **Flexible control on low cost microcontrollers**
- **Schedulability analysis for CAN-based control applications with dynamic bandwidth management**
- **Profiling and analysis**
- **Sensitivity analysis**
- **Optimal period selection and scheduling for embedded controllers**
- **Sporadic event-based Control**
- **Real-Time scheduling on multicores**
- **Dynamic run-time adaptability**
- **Integrated memory and communications management based on RTSJ**
- **Assessment of the IEEE 802.15.4 GTS scheduling and allocation mechanism for real-time WSNs**
- **Combined Energy and QoS management in WSNs**

This section was already presented in the Y2 deliverable, in section 1.7 and 3.1.

1.10 Work achieved in Year 2 (Jan-Dec 2009)

The problems tackled during Year 2 have focused on the milestones defined in the Year 1 deliverable. Here a brief description is provided; further details are to be found in the following Technical Achievements (TA) section, which are numbered for convenience.

Extend the taxonomy of resource usage

The main cluster-wide problem tackled during the last 12 months is the establishment of a wiki-based taxonomy of resource usage that has started during the first year. This wiki is now open for external read-only access. Within this taxonomy resources have been categorised and characterised. Approaches to structuring multi-resource platforms are being considered including the use of banded notions of time, hierarchical structures and a resource usage model. See TA1. Input to the taxonomy will also be taken from many of the other activities, particularly TA17.

Extend the use of hierarchical and contract-based scheduling to multi-resource systems

Hierarchical scheduling is gaining increased attention in particular in virtualised environments. The virtualisation may be chosen for composability reasons like the 653-2 standard or may be driven by considerations of security (i.e. separation of data in different partitions), robustness (e.g. avoidance of life style applications threatening the basic communication infrastructure in mobile phones) or data privacy reasons (e.g. preventing identity theft in banking applications). See TA2. Also, work undertaken by University Carlos III of Madrid.

Produce effective scheduling and placement algorithms for multiprocessor systems

Multiprocessor scheduling (including multicore) continues to be a topic that is being actively pursued in many centres world-wide, including a number in the ArtistDesign consortium. A significant number of publications have been produced in this area and a number of approaches are being developed, including the use of resource reservation techniques. See TA7, TA9 and TA15. TA16 also considers issues concerned with scheduling communications that are necessary with multiprocessor systems.

Extend sensitivity analysis to EDF and multiprocessor systems.

This topic has been pursued at York (with Professor Sanjoy Baruah of North Carolina). A set of techniques have been produced that allow sensitivity analysis to be undertaken via a single pass of a schedulability test. An internal report is available.

Produce mode change algorithms suitable for multiprocessor systems.

Global and preemptive scheduling problem of multi-mode real-time systems upon identical multiprocessor platforms has been tackled by the Polytechnic Institute of Porto. See TA11. Also, work undertaken by University Carlos III of Madrid.

Determine an effective way of undertaking (static) architectural tradeoffs.

Work is starting on this via the INDEXYS project.

Determine an effective way of undertaking (dynamic) adaptive resource management (making use of feedback techniques from the control environment).

Issues concerning control and scheduling (and resource management) have been considered in TA5, TA6 and TA10. Period selection is addressed in TA13, and event based control problems in TA14. Work has been undertaken by Carlos III University of Madrid. Issues concerned with education in this topic are covered in TA12.

Define a framework that can accommodate multiple time-frames within a single system and facilitate hierarchical scheduling, cascade control and other means of separating temporal concerns.

Work on the timeband notion has continued at York and is contributing to the taxonomy work. See TA1.

The second year deliverable provided more detail on the technical achievement (TAs) under the following headings:

- **Development of the taxonomy as a wiki**
- **Hierarchical scheduling**
- **Language support for programming schedulable systems**
- **Memory resource management**
- **Deadline-period analysis in embedded control systems**
- **Resource management for control tasks**
- **Extending resource reservation to multicore systems**
- **Limited pre-emptive scheduling**
- **Produce effective scheduling and placement algorithms for multiprocessor systems**
- **Effective way of undertaking (dynamic) adaptive resource management**
- **Mode change algorithms suitable for multiprocessor systems**
- **Virtual educational laboratory of distributed process control**

- **Optimal period selection and scheduling for embedded controllers**
- **Sporadic event-based control**
- **Task concurrency management and access scheduling for dynamic systems**
- **Techniques to avoid beacon collisions in IEEE 802.15.4 cluster-tree networks**
- **Multi-resource contract-based scheduling framework**
- **Dataflow scheduling using constraint programming**
- **Low-level adaptation mechanisms in contract-based systems**

This section was already presented in the Y2 deliverable, in sections 1.8 and 3.1.

1.11 Problems Tackled in Year 3 (Jan-Dec 2010)

The problems tackled during Year 3 have focused on the milestones defined in the Year 2 deliverable. Here a brief description is provided; further details are to be found in the following Technical Achievements (TA) section, which are numbered for convenience.

Extend the taxonomy of resource usage.

A major activity again this year has been the resource wiki that is now open for external read-only access: <http://www.cs.york.ac.uk/ArtistResourceManagement>. Within this taxonomy resources have been categorised and characterised. Approaches to structuring multi-resource platforms are being considered including the use of banded notions of time, hierarchical structures and a resource usage model. See TA1.

Extend the use of hierarchical and contract-based scheduling to multi-resource systems.

A number of projects and research activities have continued to look at means of supporting multiple applications (or components) on the same platform. These techniques need to support composition and isolation. Two important schemes that enable this to happen are hierarchical scheduling and the use of contracts. See TA2, TA3, TA7, TA9 and TA19.

Produce effective scheduling and placement algorithms for multiprocessor systems.

Multiprocessor scheduling (including multicore) continues to be a topic that is being actively pursued in many centres world-wide, including a number in the ArtistDesign consortium. A significant number of publications have been produced in this area and a number of approaches are being developed, including the use of resource reservation techniques. See TA10, TA11, TA17, TA19 and TA26. TA13, TA21 and TA22 also consider issues concerned with scheduling communications that are necessary with multiprocessor systems. Multiprocessor systems that are also distributed have been considered in TA8, TA14 and TA18.

Determine an effective way of undertaking (static) architectural tradeoffs.

Work here has focused on the emerging topic of Cyber-Physical systems and the trade-offs that must be made when energy usage has to be considered. Also considered is the way static

analysis can be used to define the slack that can be used at run-time to enhance the quality of the system's outputs. Work towards this objective is to be found in TA12, TA15, TA16 and TA20.

Determine an effective way of undertaking (dynamic) adaptive resource management (making use of feedback techniques from the control environment).

Issues concerning control and scheduling (and resource management) have been considered in TA4, TA5, TA6, TA12 and TA24.

Extend sensitivity analysis to EDF and multiprocessor systems

The work undertaken in previous years has been consolidated into a number of publications and has resulted in the development of a C=D algorithm for semi-partitioned multiprocessor scheduling. See TA26.

Support resource management software construction in programming languages

The need to be able to program systems that make use of the resource management notions developed by researchers means that support at the programming language and operating system level is required. The latter is discussed in the OS cluster deliverable. Language issues and their standardisation have been progressed in TA8 and TA25.

-- The above is new material, not present in the Y2 deliverable --

2. Summary of Activity Progress in Year 3 *(Jan-Dec 2010)*

2.1 Technical Achievements

1. Development of the Taxonomy as a wiki (all partners lead by York)

An initial definition of the taxonomy was produced in Year 2. Year 3 produced a plan for the wiki and to populate it with relevant material. This has taken time as it involved a new way of working for many people (in a very wide community). In Year 4, it will be opened up to external input. Proving bounds on the time at which an event occurs (for example the time that a command is sent to an actuator) requires that the system is modelled. In this taxonomy, we have started to model the system by expressing constraints on the values that variables (typically) can take. The variables indicate things like: how much does a specific task execute in this time interval. We have distinguished between a number of different types of constraint.

2. Development of component-based real-time systems using reactive behaviour models (Cantabria)

Research has been carried out on methods to effectively use a component-based software design process for developing real-time systems. These methods address the problem of specifying the timing behaviour properties and requirements of the individual components that may be used to compose an application. Another important aspect has been the adaptation of the deployment and configuration plan to specify the timing behaviour of the application using reactive models. As a proof of concept, all these ideas have been implemented into a set of tools and a component-based process showing that all the relevant real-time aspects of the application and its components can be specified and captured, while the tool is able to generate the code required to make the interactions between the different active components and the underlying scheduling mechanisms.

3. Real-time scheduling (Cantabria, Florida State University)

Work has been carried out with Florida State University to address some defects discovered in the POSIX specification of the sporadic server scheduling algorithm, by proposing a correction for the standard. This scheduling algorithm is very interesting for implementing resource-reservations, and useful for resource management in complex real-time systems.

Research was also carried out on methods for assigning scheduling deadlines to the threads and messages of end-to-end flows in distributed real-time systems with the goal of optimizing the time responses and the ability of the system to meet its end-to-end deadlines.

4. Optimal period selection and scheduling for embedded controllers (ULUND, Linköping Univ, UPC, SSSA)

The work on how to optimally schedule a set of distributed control algorithms studied by ULUND and Linköping University and introduced in the Y1 and Y2 deliverables has continued further and led to two joint publications. The objective is to minimize the combined cost of the controllers in the application. The design problem is solved using a genetic algorithm that decides the execution pattern of the control tasks. In the dynamic scheduling case, simulation is used to estimate the average delay and jitter of the control tasks. This year, the scheduling of so called self-triggered controllers has also been investigated, now also in collaboration with

UCLA. Similar to before the Jitterbug toolbox from Lund is used to evaluate the control performance, taking the delay and jitter into account.

At ULUND a new jitter margin result has also been developed. The new stability margin result allows tradeoffs between the amount of jitter in the sampling times of the controller and the amount of jitter in the actuation times of the controller.

The problem of optimal period and parameters assignment for controllers has also been studied by ULUND, UPC, and SSSA, which has led to two joint publications.

5. Sporadic Event-Based Control (ULUND, SSSA)

The work on event-based control carried out by ULUND continues. This year the work has focused on network scheduling of event-based controllers.

6. Dataflow Scheduling using Constraint Programming (ULUND, Ericsson)

Within the ACTORS project ULUND and Ericsson are applying constraint programming for static uniprocessor scheduling of synchronous data-flow (SDF) networks and cyclo-static data flow (CSDF) networks. A model compiler tool has been developed that automatically classifies actors written in the CAL dataflow language as either having a static input-output behaviour, i.e., being either SDF or Cyclo-Static Data Flow, or a dynamic behaviour. In the static case the tool also generates the input and output port token consumption and production rates. When the user has selected a static sub-network, the tool automatically calculates optimal schedules using constraint programming and merges the actors involved into a single actor.

7. Scheduling real-time multimedia with dynamic priorities over contract-based budgets (UC3M)

UC3M has worked on algorithms for dynamic assignment of priorities in multimedia systems that utilize contract-based scheduling and guaranteed budget assignment. This work is currently being extended in cooperation with UPM. An efficient, light-weight mechanism has been developed that makes it possible to combine budget scheduling with a fixed two-band priority scheduling mechanism to allow exhausted-budget tasks to continue execution if enough processor is available. A framework for real-time execution in multimedia systems has been extended based on HOLA-QoS that allows the integration of remote multimedia processing. In the context of iLAND (an ARTEMIS call 1 project coordinated by UC3M), this work has been extended to integrate real-time video transmission over real-time middleware as DDS. DDS priority model is being extended by UC3M.

8. Non-functional information models for distributed real-time Java (UC3M)

UC3M developed three non-functional information models for distributed real-time Java. They were evaluated on DREQUIEMI, its distributed real-time Java framework, and the empirical results of this experience have been reported to a journal and a conference. The non-functional models may improve the performance of distributed real-time Java applications significantly, and support different end-to-end performance models. The models also make it easier to integrate real-time acceptance tests and dispatching policies in distributed real-time Java infrastructures.

Also, UC3M defined a distributed cyber-architecture for distributed real-time Java which was reported in a conference paper. The architecture uses their previous model for distributed real-time Java, namely DREQUIEMI, as part of this architecture. It is based on common-off-the-shelf technology and integrates well-known technology such as RTSJ, DRTSJ, and DDS into a real-time Java centric architecture. Inside this architecture, UC3M also refined its architecture

for distributed real-time Java applications named DREQUIEMI. Their reference architecture and details on how remote invocations are carried out have also been published in a conference paper. This achievement addresses year 3 milestones: 7 and 8.

9. New time-deterministic reconfiguration algorithms for service-based real-time applications (UC3M)

As part of the efforts of developing techniques suitable for real-time reconfigurable systems, UC3M developed time-bounded reconfiguration algorithms for service-based real-time applications that take into account the real-time characteristics of the whole system.

We proposed a reconfiguration algorithm that relies on previously proposed composition algorithms that are time-bounded, and that introduces the concept of neighbourhood and a timeout to reduce the number of combinations explored. This achievement addresses year 3 milestones: 3, 7, and 8.

10. Suboptimality of Scheduling results (York, EPFL, North Carolina(US))

At York, we have previously collaborated with Thomas Rothvoss (EPFL) and Prof. Sanjoy Baruah (North Carolina) on research determining the sub-optimality of fixed priority pre-emptive scheduling (compared to an optimal pre-emptive algorithm such as EDF). This year, we extended this thread of research to non-pre-emptive systems, collaborating with Laurent George and Pierre Courbin (INRIA). We examined the relative effectiveness of fixed priority non-pre-emptive scheduling (FP-NP), compared to an optimal work conserving non-pre-emptive algorithm; Earliest Deadline First (EDF-NP). The quantitative metric used in this comparison was the processor speedup factor, defined as the factor by which processor speed needs to increase to ensure that any taskset that is schedulable according to EDF-NP can be scheduled using FP-NP. For sporadic tasksets with implicit, constrained, or arbitrary deadlines, we showed that this speedup factor is lower bounded by $1/\Omega \approx 1.76322$ and upper bounded by 2.

11. Effective hard real-time multiprocessor scheduling algorithms with low Overheads (York)

Finding provably effective, but low overhead hard real-time multiprocessor scheduling algorithms is currently a key research goal. There are two promising approaches based on (i) partitioned and (ii) global scheduling. Both approaches have been explored. Semi-partitioned algorithms allocate most tasks to a single processor, splitting just a few tasks which are the only ones able to migrate from one processor to another at run time. By contrast, minimally dynamic global scheduling algorithms permit all tasks to migrate, but limit the number of changes in the priority of each job to one.

Research work has been carried out into Fixed Priority until Zero Laxity (FPZL) scheduling. FPZL is similar to global fixed priority pre-emptive scheduling; however, whenever a task reaches a state of zero laxity it is given the highest priority. FPZL is a minimally dynamic algorithm, in that the priority of a job can change at most once during its execution, tightly bounding the number of pre-emptions. We developed polynomial time and pseudo-polynomial time sufficient schedulability tests for FPZL. These tests we subsequently improved by computing upper bounds on the amount of execution that each task can perform in the zero laxity state. Our empirical evaluation showed that FPZL is highly effective, with a significantly larger number of tasksets deemed schedulable by the tests for this algorithm, than by state-of-the-art schedulability tests for global EDF or Earliest Deadline until Zero Laxity (EDZL) scheduling.

12. Slack management for control tasks in monoprocessor systems (UPC, UCSC-Scott Brandt)

Effective slack management, i.e. management of unused computing resources, for real-time control tasks mandates to redistribute the available resources between controllers at run-time. Slack can be allocated to control tasks to alter their rate of progress via for example, the controllers' period, in order to adapt their behaviour to changes in the computing platform and in the environment. An experimental evaluation on the Erika real-time kernel of two representative slack redistribution policies has been performed. One policy represents the feedback scheduling approach, where a resource manager is responsible for modifying the rate of progress of each control task. Alternatively, the second policy represents an event-driven approach, where each control task decides its rate of progress. The evaluation indicates that the feedback scheduling approach provides better control performance while the event-driven approach is good at saving computing resources.

13. A reconfigurable hierarchical scheduling framework (Aveiro, UnivPorto and Malardalen)

Implementation and analysis of a reconfigurable hierarchical scheduling framework within an enhanced Ethernet switch was performed. This work is an evolution of the work on an FTT-enabled switch in which the asynchronous window can be divided and further subdivided in partitions that are flexible enough to allow an efficient use of bandwidth but which enforce mutual temporal and spatial isolation.

14. Utilisation-based Schedulability tests (UnivPorto, CMU and Aveiro)

Development, characterization and assessment of utilization-based schedulability tests accounting for release jitter for use in dynamic QoS management in distributed systems. Three different tests were developed and compared among them and with another existing test, establishing the limits for their applicability.

15. Investigation on a new class of Real-Time Systems: Real-Time Physical Systems (Pavia)

The work was to define the characteristics and properties of Real-Time Physical Systems (RTPS), and their application to real test cases. The relationship among RTPS and Cyber-Physical Systems, Cyber-Physical Energy Systems and Hybrid Systems has been investigated. Moreover, the bridge between RTPS, defined in terms of Real-Time Physical Resources, has been drawn to existing system classes in current real-time systems research, i.e., power-aware and temperature-aware systems. This work has been carried out at the University of Pavia. The potential and the applicability of RTPS has been verified by implementing a control system that was able to optimize the electric power demand in a set of refrigerators, while achieving the system timing and physical constraints. Experimentation was done in collaboration with an SME.

16. Power-constrained embedded system design (Pavia)

This research effort has been dedicated to the study of power-aware techniques for embedded systems when multiple sensors are required to be sampled, with different timings. The research on power-constrained embedded systems has led to the implementation of an Inertial

Measurement Unit (IMU) device, exploiting the underlying Real-Time Operating System for managing real-time constraints on sensor data sampling.

17. Design space exploration framework for run-time resource management on multi-core systems (IMEC, Politecnico de Milano and University of Lugano)

In the context of this collaboration, an automatic design space exploration (DSE) framework was developed for enabling resource management through software re-configuration on an industrial multi-core platform. At design-time DSE is used with multiple simulators running at multiple abstraction levels to quickly derive the Pareto set of operation points. At run-time a hierarchical Runtime Resource Manager (RRM) selects the optimal operation point while introducing a marginal overhead.

18. Software Transactional Memory for Real-Time Systems (Porto, York)

The foreseen evolution of chip architectures to higher numbers of heterogeneous cores, with non-uniform memory and non-coherent caches brings renewed attention to the use of Software Transactional Memory (STM) as an alternative to lock-based synchronisation. However, STM relies on the possibility of aborting conflicting transactions to maintain data consistency, which impacts on the responsiveness and timing guarantees required by real-time systems. In these systems, contention delays must be (efficiently) limited so that the response times of tasks executing transactions are upper bounded and task sets can be feasibly scheduled. In this work, we are analyzing the role of the transaction contention manager to reduce the number of transaction retries and to help the real-time scheduler, assuring schedulability. For such purposes, the contention management policy must be aware of on-line scheduling information. The research has resulted in a preliminary paper at the RTSS 2010 WiP session <http://tcrts.org/rtss-2010/wip/>

19. Parallelism in Server-based Multiprocessors (Porto)

The constant-bandwidth server abstraction has proved very useful in designing, implementing, and reasoning about single core open real-time systems. However, modern open real-time systems increasingly generate heavy workloads and it is rapidly becoming unreasonable to expect to implement them as single core systems. In fact, a general shift from uncore to multicore processors can be seen both in the general purpose and embedded domains as an energy-efficient way to boost application performance. Therefore, there have been significant efforts to extend reservation-based real-time scheduling theory to make it applicable to multiprocessor systems as well. Nevertheless, all these works consider task models where tasks use at most a single core at each time instant. This work analyses the increased need to support dynamic task-level parallelism in open real-time systems and proposes the general guidelines of a novel scheduling approach that combines a work-stealing load balancing policy with a multicore reservation-based approach. This research has resulted in a preliminary paper at the RTSS 2010 WiP session <http://tcrts.org/rtss-2010/wip/>

20. Effective Energy Management in Real-Time Systems (Porto)

Energy Management in Systems is becoming an increasing concern due to the fast rising number of battery operated devices. At the same time, the reduced size of individual transistors in hardware design reduces the ratio of dynamic vs. static power consumption and subsequently reduces the attractiveness of dynamic voltage and frequency scaling (DVFS) and increases that of employing processor sleep states to obtain energy savings. Existing work on energy management for real-time systems often ignores the substantial time and energy costs

of making DVFS and sleep state transitions, and/or assumes very simple models. Researchers at CISTER have developed a comprehensive technique to take the cost of those transitions into account while maintaining all real-time guarantees.

21. Scheduling algorithms for real-time and reliable communication in IEEE 802.11 networks (Catania)

This activity dealt with a scheduling framework and a scheduling algorithm to handle real-time communication over the wireless section of a hybrid wired/wireless industrial network, which uses a scheduled protocol that extends the IEEE 802.11n standard. The proposed scheduling framework is able to schedule real-time flows traversing one or multiple Access Points, providing both intra- and inter-cell scheduling. Moreover, the scheduling framework offers support for three different kinds of real-time flows; periodic, aperiodic, and sporadic. A simulative analysis is provided to assess the computation times as a function of the number of flows and the number of schedulable flows in different scenarios.

22. Load management in flexible industrial networks characterized by a hybrid wired/wireless architecture and mobile wireless nodes (Catania)

A novel load balancing algorithm that takes into account both the bandwidth utilization and the state of the wireless channel, so as to compensate dynamically for fluctuations in the wireless link characteristics as well as the load of mobile nodes that join and leave the AP, is proposed. The approach is based on migration of nodes from an overloaded AP to another AP capable of supporting their traffic. Simulation results highlight the advantage of the proposed technique over both the standard AP selection mechanism and previously known load balancing algorithms based on the network throughput.

23. Frequency hopping management mechanism for Bluetooth networks used in industrial environments (Catania)

In large industrial networks, multiple piconets can be deployed to work in parallel within a limited area. If they overlap, it is likely that a node transmission in a piconet will be affected by the co-channel interference due to the transmission, on the same channel and at the same time, of a node belonging to any of the co-located piconets.

A novel frequency hopping management mechanism for Bluetooth networks used in industrial environments, called a Piconet Management Approach, has been proposed. The PMA allows a controller node to synchronize the clocks of the overlapping piconets and manage their frequencies, i.e. hopping sequences, to limit the interference between adjacent cells. This is accomplished by introducing a variant in the standard frequency hopping algorithm which does not contrast with the operational modes defined in the BT standard.

The results obtained through ns-2 simulations showed that the PMA achieves significant improvements in terms of collision reduction, thus leading to better network throughput, reduced packet loss and delay.

24. Dynamic resource management (Universidad Politécnica de Madrid. UPM)

UPM developed a component for dynamic resource management that can be adapted to different kernels. UPM has extended this implementation to support multi-core architectures and improved its performance.

There are a wide range of applications that must be written using non-real-time languages for a variety of reasons, such as library availability or the need to use legacy code. However, they may also have time requirements. The purpose of this work has been to provide some support

for improving the predictability of activities with this type of requirement. UPM has worked on providing budgets or resource reserves to specific threads for this purpose.

25. Language support for Programming Schedulable Systems (York, Cantabria, Porto, Madrid, Valencia)

Standardisation work continues for the programming languages Ada and Java and how they can be used to program systems that have timing constraints and must therefore control the management of resources. This year the work has focused on getting support for multiprocessors into the next versions of Ada and the Real-Time Specification for Java. These have now effectively been agreed and will enter into the standards at their next releases.

26. Sensitivity analysis and task splitting (York and UNC)

The previous work on sensitivity analysis for single processor EDF scheduling systems has been completed and published. An extension for multiprocessor systems has been developed that allows a single task per processor to be split into two parts. The first part has the constraint that its computation time must be equal to its deadline (hence the name of the algorithm $C=D$), the second part has a normal behaviour. Early results imply that this is a very effective scheme both theoretically and in practice.

27. Resource management (TUKL)

TUKL have conducted research into abstraction layers and resource allocation mechanisms for real-time resource management in order to decrease quality of service degradation upon scarce resource availability. As an application example, we have developed an adaptive MPEG-2/MPEG-4 video streaming demo.

28. Scheduling algorithms for adaptive and flexible real-time systems (TUKL)

TUKL have conducted research into adaptive and flexible real-time scheduling for improved resource usage. This work includes scheduling of target sensitive applications with the gravitational task model, and integration of off-line and on-line scheduling based on the slot shifting algorithm.

29. Hierarchical scheduling in hypervisors systems (UP Valencia)

This activity is focused on the design and implementation of hierarchical scheduling policies in hypervisor-based environments where a virtualisation layer provides virtual machines to execute partitions that include guest operating systems. Guest operating systems are real-time operating systems para-virtualised to be executed as a virtual machine.

Novel design techniques have been developed for the cyclic scheduler at the hypervisor level taking into account the temporal requirements of the partition's tasks. The task model is based on end to end flow activities and the result of the analysis is cyclic scheduling that considers all the temporal constraints and precedence relations between tasks. One of the differential aspects is the approach to the virtual services costs: partition context switch, communication between partitions, synchronization between the temporal execution windows of a partition and the internal tasks and security issues related to the confidentiality related to shared resources.

XtratuM is the hypervisor developed by UP Valencia that is used as an experimentation platform to evaluate these developments. This hypervisor is being used by the aerospace industry to develop temporally and spatially partitioned systems.

-- The above is new material, not present in the Y2 deliverable --

2.2 Individual Publications Resulting from these Achievements

University of York

A. Burns and B. Littlewood, Reasoning about the Reliability of Multi-Version, Diverse Real-Time Systems, Proceeding of the 31st Real-Time Systems Symposium, pp73-81, 2010.

Alan Burns Ian J. Hayes, A Timeband Framework for Modelling Real-Time Systems, Real-Time Systems, Vol 45, pp106-142, 2010.

Kun Wei, Jim Woodcock, Alan Burns, A Timed Model of Circus with the Reactive Design Miracle, 8th International Conference on Software Engineering and Formal Methods (SEFM), pp 315-319, 2010.

Alexandros Zerzelidis, Andy Wellings, A framework for flexible scheduling in the RTSJ, ACM Trans. Embed. Comput. Syst., Vol 10, pp1-44, 2010.

Abdul Haseeb Malik, Andy Wellings, Yang Chang, A locality model for the real-time specification for Java, JTRES '10: Proceedings of the 8th International Workshop on Java Technologies for Real-Time and Embedded Systems, pp36-45, 2010.

A.J. Wellings, A.H. Malik, N.C. Audsley, A. Burns, Ada and cc-NUMA architectures what can be achieved with Ada 2005?, Ada Letters, Vol 30, pp125-134, 2010.

Andy Wellings, MinSeong Kim, Asynchronous event handling and safety critical Java, JTRES '10: Proceedings of the 8th International Workshop on Java Technologies for Real-Time and Embedded Systems, pp53-62, 2010.

A. Burns, A.J. Wellings, Dispatching Domains for Multiprocessor Platforms and their Representation in Ada, Proceedings of Reliable Software Technologies - Ada-Europe, pp41-53, 2010.

Fengxiang Zhang, Alan Burns, Dividing point value selections for Improved Quick Processor-demand Analysis, Software Technology and Engineering (ICSTE), pp170-175, 2010.

A. Burns, Dual Priority Scheduling: Is the Processor Utilisation bound 100%?, Proceedings of 1st International Real-Time Scheduling Open Problems Seminar, RTSOPS 2010, pp3-4.

MinSeong Kim, Andy Wellings, Efficient asynchronous event handling in the real-time specification for Java, ACM Trans. Embed. Comput. Syst., Vol 10, pp 1-34, 2010.

A.J. Wellings, A. Burns, Generalizing the EDF scheduling support in Ada 2005, Ada Letters, Vol 30, pp116-124, 2010.

R.I. Davis, A. Burns, Improved priority assignment for global fixed priority pre-emptive scheduling in multiprocessor real-time systems, Real-Time Systems, Vol 40, pp1-40, 2010.

Jim Woodcock, Marcel Oliveira, Alan Burns, Kun Wei, Modelling and Implementing Complex Systems with Timebands, Secure System Integration and Reliability Improvement, pp1-13, 2010

A. Burns, R.I. Davis, P. Wang, F. Zhang, Partitioned EDF Scheduling for Multiprocessors using a C=D Scheme, Proceedings of 18th International Conference on Real-Time and Network Systems (RTNS), pp169-178, 2010.

R.I. Davis, L. George, P. Courbin, Quantifying the Sub-optimality of Uniprocessor Fixed Priority Non-Pre-emptive Scheduling, Proceedings of 18th International Conference on Real-Time and Network Systems (RTNS), pp1-10, 2010.

Y. Chang, R.I. Davis, A.J. Wellings, Reducing Queue Lock Pessimism in Multiprocessor Schedulability Analysis, Proceedings of 18th International Conference on Real-Time and Network Systems (RTNS), pp99-108, 2010.

Z. Shi, A. Burns, L. S. Indrusiak, Schedulability Analysis for Real Time On-Chip Communication with Wormhole Switching, Int Journal of Embedded and Real-Time Communication Systems (IJERTCS), pp1-22, 2010.

A. Burns, C. Dale, Scheduling and Timing Analysis for Safety-Critical Real-Time Systems, Electronics World, Vol 116, pp18-20, 2010.

F. Zhang, A. Burns, S. Baruah, Sensitivity Analysis for EDF Scheduled Arbitrary Deadline Real-Time Systems, Proceedings of 16th IEEE Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA), pp61-70, 2010.

F. Zhang, A. Burns, S. Baruah, Sensitivity analysis of relative deadline for EDF scheduled real-time systems, Proceedings of 2nd International Conference on Mechanical and Electronics Engineering (ICMEE), pp296-301, 2010.

Fengxiang Zhang, Alan Burns, and Sanjoy Baruah, Sensitivity Analysis of the Minimum Task Period for Arbitrary Deadline Real-Time Systems, Proceedings of the 16th IEEE Pacific Rim International Symposium on Dependable Computing (PRDC), Tokyo, Japan. pp 101-108, 2010.

F. Zhang, A. Burns, S. Baruah, Sensitivity analysis of task period for EDF scheduled arbitrary deadline real-time systems, Proceedings of 3rd IEEE International Conference on Computer Science and Information Technology (ICCSIT), pp23-28, 2010.

A. Burns, A.J. Wellings, Supporting Execution on Multiprocessor Platforms, Proceedings of 14th International Real-Time Ada Workshop, Vol XXX, IRTAW, 2010.

Niusha Hakimipour, Paul Strooper, Andy Wellings, TART: Timed-Automata to Real-Time Java Tool, 8th International Conference on Software Engineering and Formal Methods (SEFM), pp299-309, 2010.

F. Zhang, A. Burns, S. Baruah, Task parameter computations for constraint deadline real-time systems with EDF scheduling Proceedings of International Conference on Computer Design and Applications (ICCD), 2010.

A.J. Wellings, A. Burns, The Evolution of Real-Time Programming Revisited: Programming the Giotto Model in Ada 2005, Proceedings of Reliable Software Technologies - Ada-Europe, pp196-207, 2010.

A.J. Wellings, A. Burns, User-defined clocks is it the right time now?, Ada Letters, Vol 30, pp104-115, 2010.

MinSeong Kim, Andy Wellings, Using the executor framework to implement asynchronous event handling in the RTSJ, JTRES '10: Proceedings of the 8th International Workshop on Java Technologies for Real-Time and Embedded Systems, pp16-25, 2010.

Scuola Superiore Sant'Anna, Pisa

G. Lipari and E. Bini, A Framework for Hierarchical Scheduling on Mutiprocessors: From Application Requirements to Run-Time Allocation, Proc IEEE Real-Time System Symposium, pp 249-258, 2010.

Gang Yao and Giorgio Buttazzo, "[Reducing Stack with intra-task Threshold Priorities in Real-Time Systems](#)", Proceedings of the 10th International Conference on Embedded Software (EMSOFT 2010), Scottsdale, Arizona (USA), October 24-29, 2010.

L. Santinelli, M. Marinoni, F. Prosperi, F. Esposito, G. Franchino, and G. Buttazzo, "[Energy-Aware Packet and Task Co-Scheduling for Embedded Systems](#)", Proceedings of the 10th International Conference on Embedded Software (EMSOFT 2010), Scottsdale, Arizona (USA), October 24-29, 2010.

Gianluca Franchino, Giorgio Buttazzo and Mauro Marinoni, "[An Energy-Aware Algorithm for TDMA MAC Protocols in Real-Time Wireless Networks](#)", Proceedings of the IEEE Symposium on Industrial Embedded Systems, (SIES 2010), Trento, Italy, July 7-9, 2010.

Gang Yao, Giorgio Buttazzo, Marko Bertogna, "[Comparative evaluation of limited preemptive methods](#)", Proceedings of the 15th IEEE International Conference on Emerging Technology and Factory Automation (ETFA10), Bilbao, Spain, September 13-16, 2010.

Gang Yao, Giorgio Buttazzo and Marko Bertogna, "[Feasibility Analysis under Fixed Priority Scheduling with Fixed Preemption Points](#)", Proceedings of the 16th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA 2010), Macau, China, August 23-25, 2010 (Best Paper).

Giorgio Buttazzo, Enrico Bini, and Yifan Wu, "[Partitioning Parallel Applications on Multiprocessor Reservations](#)", Proceedings of the 22nd Euromicro Conference on Real-Time Systems (ECRTS 10), Brussels, Belgium, July 6-9, 2010.

Marko Bertogna, Giorgio Buttazzo, Mauro Marinoni, Gang Yao, Francesco Esposito, and Marco Caccamo, "[Preemption Point Placement for Sporadic Task Sets](#)", Proceedings of the 22nd Euromicro Conference on Real-Time Systems (ECRTS 10), Brussels, Belgium, July 6-9, 2010.

Giorgio Buttazzo, Enrico Bini, and Yifan Wu, "[Heuristics for Partitioning Parallel Applications on Virtual Multiprocessors](#)", Proceedings of the First International Workshop on Adaptive Resource Management (WARM 2010), Stockholm, Sweden, April 12, 2010.

Universidad de Cantabria

Laura Barros, Ángela del Barrio, Patricia López Martínez, and José M. Drake. "New container services for the integration of Component-based Applications on Complex Industrial Platforms". 15th IEEE International Conference on Emerging Technologies and Factory Automation, ETFA 2010, Bilbao, September 2010, IEEE,

Patricia López Martínez and José M. Drake (Thesis Supervisor). Phd Thesis: "Development of component-based real-time systems using reactive behaviour models" (In Spanish), Phd. Thesis, University of Cantabria, September 2010. <http://www.tesisenred.net/TDR-0928110-111910>.

Juan M. Rivas, J. Javier Gutiérrez, J. Carlos Palencia, and Michael González Harbour. "Optimized Deadline Assignment and Schedulability Analysis for Distributed Real-Time Systems with Local EDF Scheduling" 8th International Conference on Embedded Systems and Applications, ESA'2010, Las Vegas (Nevada, USA), July 2010.

Patricia López Martínez, Laura Barros and José M. Drake. "Scheduling Configuration of Real-Time Component-Based Applications". 15th Int. Conf. On Reliable Software Technologies, Ada-Europe'2010, Valencia (Spain), in Lecture Notes in Computer Science, LNCS Vol. 6106, pp. 181-195, June 2010.

Mark Stanovich, Theodore Baker, Andy Wang and Michael Gonzalez Harbour. "Defects of the POSIX Sporadic Server and How to Correct Them". 16th IEEE Real-Time and Embedded Technology and Applications Symposium. Stockholm, Sweden, April, 2010.

TUCL

Raphael Guerra, Gerhard Fohler, An optimum generalized equilibrium solution for the gravitational task model, RTNS '10: Proceedings of the 18th International Conference on Real Time and Network Systems, IRIT Lab, Toulouse, France, November 2010.

Stefan Schorr, Gerhard Fohler, Online Admission of Non-Preemptive Aperiodic Tasks in Offline Schedules, Proceedings of Work-in-Progress Session, 22nd Euromicro Conference on Real-Time Systems 2010, ECRTS, Brussels, Belgium, Work in Progress Session, July 2010.

Anand Kotra, Gerhard Fohler, Demo : Resource aware real-time stream adaptation for MPEG-2 transport streams in constrained bandwidth networks, Proceedings of IEEE International Conference on Multimedia & Expo (ICME), singapore, July 2010.

Rodrigo Ferreira Coelho, Anand Kotra, Gerhard Fohler, A Control Theory Approach to Video Stream Adaptation for Restricted Bandwidth Networks, Proceedings of 1st Workshop on Adaptive Resource Management (WARM10), Stockholm, Sweden, April 2010.

Lund University

T. Henningsson and A. Cervin, A Simple Model for the Interference Between Event-Based Control Loops Using a Shared Medium, In Proc. 49th IEEE Conference on Decision and Control, Atlanta, GA, December 2010.

Polytechnic Institute of Porto

Björn Andersson, Luis Miguel Pinho, "Implementing Multicore Real-Time Scheduling Algorithms Based on Task Splitting Using Ada 2012", in Proceedings of the 15th International Conference on Reliable Software Technologies (Ada-Europe 2010), Valencia, Spain.

Stefan M. Petters, Muhammad Ali Awan, "Slow Down or Race to Halt: Towards Managing Complexity of Real-Time Energy Management Decisions", Proceedings of the 12th Brazilian Workshop on Real-Time and Embedded Systems (WTR 2010), Work-in-Progress Session, 24.05.2010, Gramado/RS, Brazil.

António Barros, Luís Miguel Pinho, "Managing contention of software transactional memory in real-time systems", Real-Time Systems Symposium, Work-in-Progress session, 2010.

Luis Nogueira, Luís Miguel Pinho, "Supporting Parallelism in Server-based Multiprocessor Systems", Real-Time Systems Symposium, Work-in-Progress session, 2010.

IMEC

Couvreur, C.: Exploration framework for run-time resource management of embedded multi-core platforms. In IEEE Int. Conf. on Embedded Computer Systems: Architectures, Modeling, and Simulation, Samos, Greece (2010)

Silvano, C.; Fornaciari, W.; Crespi Reghizzi, S.; Agosta, G.; Palermo, G.; Zaccaria, V.; Bellasi, P.; Castro, F.; Corbetta, S.; Di Biagio, A.; Tartara, M.; Siorpaes, D.; Hubert, H.; Stabernack, B.; Brandenburg, J.; Palkovic, M.; Raghavan, P.; Couvreur, C.; Bartzas, A.; Xydis, S.; Soudris, D.; Kempf, T.; Ascheid, G.; Meyr, H.; Ansari, J.; Mahonen, P. and Vanthournout, B.: 2PARMA: Parallel paradigms and run-time management techniques for many-core architectures. In IEEE Computer Society Annual Symposium on VLSI – ISVLSI, Lixouri, Greece, (2010)

Silvano, C.; Fornaciari, W.; Palermo, G.; Zaccaria, V.; Castro, F.; Martinez, M.; Bocchio, S.; Zafalon, R.; Avasare, P.; Vanmeerbeeck, G.; Couvreur, C.; Wouters, M.; Kavka, C.; Onesti, L.; Turco, A.; Bondi, U.; Mariani, G.; Posadas, H.; Villar, E.; Wu, C.; Dongrui, F.; Hao, Z. and Shibin, T.: MULTICUBE: Multi-objective design space exploration of multi-core architectures. In *IEEE Computer Society Annual Symposium on VLSI – ISVLSI*, Lixouri, Greece, (2010)

Mariani, G.; Avasare, P.; Vanmeerbeeck, G.; Couvreur, C.; Palermo, G.; Silvano, C. and Zaccaria, V.: An industrial design space exploration framework for supporting run-time resource management on multi-core systems. In *Proceedings Design, Automation and Test in Europe Conference – DATE*, Dresden, Germany (2010)

Avasare, P.; Couvreur, C.; Vanmeerbeeck, G.; Mariani, G.; Palermo, G.; Zaccaria, V. and Silvano, C.: Linking run-time management with design space exploration at multiple abstraction levels. In *DATE Workshop: "Designing for Embedded Parallel Computing Platforms"*, Dresden, Germany (2010)

Universidad Politécnica de Valencia

I. Ripoll, A. Crespo, Tutorial: Hypervisor Technology for Building Safety-Critical Systems: XtratuM. 16th International Conference on Reliable Software Technologies – Ada-Europe 2011.Valencia. June 2010.

A. Crespo, I. Ripoll, M. Masmano: Partitioned Embedded Architecture Based on Hypervisor: The XtratuM Approach. Eighth European Dependable Computing Conference, pp: 67-72, Valencia, Spain, 28-30 April 2010

C. Jouvray, G. Chartier, N. François, I. Ripoll, M. Masmano, A. Crespo: Enforcing trust in control automotive platforms. EDCC-CARS 2010: 43-46.

S. Saez, A. Crespo: Preliminary Multiprocessor Support of Ada 2012 in GNU/Linux Systems. Ada-Europe 2010: 68-82

V. Brocal, M. Masmano, I. Ripoll, A. Crespo, P. Balbastre, J.J. Metge. Xoncrete: a Scheduling Tool for Partitioned Real-Time Systems. Embedded Real Time Software and Systems (ERTS2 2010). May 2010. Toulouse (France). 2010

M. Masmano, I. Ripoll, A. Crespo, S. Peiro. XtratuM for LEON3: an OpenSource Hypervisor for High-Integrity Systems. Embedded Real Time Software and Systems (ERTS2 2010). May 2010. Toulouse (France). 2010

I. Ripoll, A. Crespo, M. Masmano, V. Brocal, P. Balbastre, S. Peiró, P. Arberet, J.J. Metge

Configuration and Scheduling tools for TSP systems based on XtratuM. In DASIA 2010. DAta Systems In Aerospace. May. Budapest 2010.

G. Herrgott, T. Pareaud, L. Planche, V. Kollias, N. Pogkas, A. Crespo, I. Ripoll, M. Masmano. Securely Partitioning Spacecraft Computing Resources. In DASIA 2010. DAta Systems In Aerospace. May. Budapest 2010.

M. Masmano, I. Ripoll, Y. Valiente, P. Balbastre and A. Crespo. LithOS: a ARINC 653 guest operating for XtratuM. Twelfth Real-Time Linux Workshop. October 2010. Nairobi, Kenya. 2010

Universidad Politécnica de Madrid

Alejandro Alonso, Emilio Salazar, Jorge López, *Resource Management for Enhancing Predictability in Systems with Limited Processing Capabilities*, 15th IEEE ETFA 2010 (Emerging Technologies and Factory Automation). 13-16 September, Bilbao (SPAIN)

University Carlos III of Madrid

M. García-Valls, Iria Estévez-Ayres, Pablo Basanta-Val. *Dynamic priority assignment scheme for contract-based QoS resource Management*. 7th IEEE International Conference on Embedded Software and Systems. (IEEE ICESSE 2010). IEEE Computer Society Press. Bradford, UK. May 29 - June 1, 2010.

M. García-Valls, Pablo Basanta-Val, Iria Estévez-Ayres. *A component model for homogeneous implementation of reconfigurable service-based distributed real-time applications*. In International Workshop on Distributed Architecture modelling for Novel component based Embedded Systems, DANCE 2010. IEEE Computer Society Press. Tozeur, Tunisia. May 30, 2010.

M. García-Valls, L. Fernández-Villar, I. Rodríguez-López, I. Estévez-Ayres, and P. Basanta-Val. *Towards a middleware architecture for deterministic reconfiguration of service-based networked applications*. 15th IEEE ETFA 2010 (Emerging Technologies and Factory Automation). 13-16 September, Bilbao (SPAIN) pp 1-8. 2010

P. Basanta-Val, M. García-Valls and I. Estévez-Ayres. Fine tuning of the multiplexing facilities of Java's Remote Method Invocation. Accepted for Concurrency and Computation Practice and Experience.

P. Basanta-Val, M. García-Valls and I. Estévez-Ayres. Extending the Concurrency Model of the Real-Time Specification for Java. Accepted for Concurrency and Computation: Practice and Experience.

P. Basanta-Val, M. García-Valls, I. Estévez-Ayres. A Neutral Architecture for Distributed Real-Time Java based on RMI and RTSJ. ETFA 2010 (Emerging Technologies and Factory Automation). 13-16 September, Bilbao (SPAIN) pp 1-8. 2010.

P. Basanta-Val, M. García-Valls, and I. Estévez-Ayres. Towards propagation of non-functional information in distributed real-time Java. IEEE ISORC 2010 (May, 5-6 2010), pp 225-232, 2010.

P. Basanta-Val, M. García-Valls, and I. Estevez-Ayres. No-Heap Remote Objects for Distributed Real-time Java. ACM Trans. in Embedded Computing Systems. Volume 10 Issue 1, August 2010.

P. Basanta-Val, I. Estévez Ayres, M. García-Valls and Luís Almeida. A synchronous scheduling service for distributed real-time Java. IEEE Trans. on Parallel and Distributed Systems . Volume 21 Issue 4 April 2010 pp. 506-520.

P Basanta-Val, M. García-Valls and I. Estévez-Ayres. Towards a Cyber-Physical Architecture for Industrial Systems via Real-Time Java Technology. International Symposium on Advanced Topics on Embedded Systems and Applications (ESA2010) Bradford (UK). June 29 - June - 1 July, 2010. pp. 2341-2346, 2010.

I. Estévez-Ayres, M.García-Valls and P. Basanta-Val. On the reconfiguration of service-based real-time applications. In Proc. Of WiP session of IEEE RTAS-2010, Stockholm, Sweden, 12-15 April, 2010.

J. Díez-Sánchez (advisor: I. Estévez-Ayres). Implementation and evaluation of composition algorithms for distributed real-time service-based applications. Master Thesis, Universidad Carlos III de Madrid. March 2010

University of Catania

M. Collotta, L. Lo Bello, O. Mirabella, An innovative frequency hopping management mechanism for Bluetooth-based industrial networks, in the proceedings of the IEEE 2010 IEEE International Symposium on Industrial Embedded Systems (SIES'10), pp. 45 – 50, 7-9 July 2010, Trento, Italy, ISBN: 978-1-4244-5839-4.

E. Toscano, F. Misenti, L. Lo Bello, "A traffic scheduler for real-time wireless communication in adaptable industrial automation systems," IEEE International Conference on Emerging Technologies and Factory Automation (ETFA), Spain, Sept. 2010, ISBN: 978-1-4244-6849-2.

M. Collotta, L. Lo Bello, E. Toscano, O. Mirabella, "Dynamic load balancing techniques for flexible wireless industrial networks," in Proc. of the 36th Annual Conference of the IEEE Industrial Electronics Society, IECON 2010, 7-10 Nov. 2010, Phoenix, US.

University of Pavia

Marco Della Vedova, Michele Ruggeri, and Tullio Facchinetti, "On real-time physical systems", in Proceedings of the 18th International Conference on Real-Time and Network Systems (RTNS), Toulouse, France, pp. 41-49, November, 2010. ISBN: 978-2-917490-11-2.

Tullio Facchinetti and Marco L. Della Vedova, "Real-Time Modeling and Control of a Cyber-Physical Energy System", in Proceedings of the First International Workshop on Energy Aware Design and Analysis of Cyber Physical Systems (WEA-CPS), Stockholm, Sweden, April, 2010.

Tullio Facchinetti, Alberto Savioli, and Emanuele Goldoni, "Design and Development of a Real-Time Embedded Inertial Measurement Unit", in Proceedings of the 25th ACM Symposium on Applied Computing (SAC'10), Sierre, Switzerland, March, 2010.

Tullio Facchinetti and Matteo Foppiano, "Resource partitioning for real-time processing on a multicore architecture", in Proceedings of the 2010 ACM Symposium on Applied Computing (SAC'10), Sierre, Switzerland, pp. 359-360, 2010.

-- The above are new references, not present in the Y2 deliverable --

2.3 Interaction and Building Excellence between Partners

The partners involved in this activity are part of an international community that meets regularly at the main conferences for the disciplines covering scheduling and resource management. These conferences include the IEEE International symposia RTSS and RTAS, and the European conference on real-time issues, ECRTS. Interactions also take place as part of other funded EU projects such as ACTORS. Finally ArtistDesign specific meetings have taken place, for example a meeting in Brussels and Kaiserslautern. Other interactions include:

- ULUND, SSSA, TUKL, EPFL, Ericsson, and Evidence continue their collaboration on feedback-based resource scheduling for multimedia terminals within the EU STREP project ACTORS (Adaptivity and Control of Resources in Embedded Systems).
- ULUND and Linköping University collaborate on integrated scheduling and synthesis of distributed and embedded control applications within the Swedish ELLIIT project.

- UC3M and Pisa have started collaboration in two areas: QoS-based resource management for embedded multimedia platforms, and service oriented architectures. Currently, UC3M is hosting a PhD student from Pisa for a one -year stay to collaborate on these topics.
- UC3M and UPM are collaborating to enhance the dynamic priority assignment protocol initially to be validated for multimedia systems and to study its validation on hard real-time platforms.
- Researchers from the University of York have proposed new extensions to the Ada programming language for real-time scheduling on multiprocessors. CISTER researchers in Porto have studied how these language extensions can be used to implement a new class of real-time scheduling algorithms called task-splitting algorithms. The results were presented in a conference paper at the conference Ada-Europe 2010.
- A student from Aveiro and UnivPorto spent 5 months at Malardalen to work on the analysis of the hierarchical scheduling framework that was implemented in a specialized Ethernet switch.
- Two students from TUKL participated in a summer school in Autrans, France, September 5-10, 2010. The summer school consisted of several workshops on cyber-physical systems.

-- The above are all Year 3 specific activities

2.4 Joint Publications Resulting from these Achievements

Yifan Wu, Giorgio Buttazzo, Enrico Bini, and Anton Cervin, "Parameter Selection for Real-time Controllers in Resource-Constrained Systems", IEEE Transactions on Industrial Informatics, to appear, 2010.

Nikolay Stoimenov, Lothar Thiele, Luca Santinelli, and Giorgio Buttazzo, "Resource Adaptations with Servers for Hard Real-Time Systems", Proceedings of the 10th International Conference on Embedded Software (EMSOFT 2010), Scottsdale, Arizona (USA), October 24-29, 2010.

Soheil Samii, Petru Eles, Zebo Peng, Paulo Tabuada, Anton Cervin, Dynamic Scheduling and Control-Quality Optimization of Self-Triggered Control Applications, In Proc. 31st IEEE Real-Time Systems Symposium, San Diego, CA, December 2010.

Soheil Samii, Anton Cervin, Petru Eles, Zebo Peng, Runtime Trade-Offs Between Control Performance and Resource Usage in Embedded Self-Triggered Control Systems, In Proc. Workshop on Adaptive Resource Management (WARM 2010), Stockholm, Sweden, April 2010.

Yifan Wu, Giorgio Buttazzo, Enrico Bini, Anton Cervin, Parameter Selection for Real-Time Controllers in Resource-Constrained Systems, IEEE Transactions on Industrial Informatics, 2010. Nov 2010 pg. 610-620

Anton Cervin, Manel Velasco, Pau Martí, Antonio Camacho, Optimal On-Line Sampling Period Assignment: Theory and Experiments, IEEE Transactions on Control Systems Technology, Issue 99, pg. 1-9, July 2010.

M. Velasco, P. Martí, J.M. Fuertes, C. Lozoya, S. Brandt. Experimental Evaluation of Slack Management in Real-Time Control Systems: Coordinated vs. Self-Triggered Approach. In Journal of Systems Architecture, Vol. 56, N. 1, January 2010.

Tullio Facchinetti, Enrico Bini, and Marko Bertogna, "Reducing the Peak Power through Real-Time Scheduling Techniques in Cyber-Physical Energy Systems", in Proceedings of the First International Workshop on Energy Aware Design and Analysis of Cyber Physical Systems (WEA-CPS), Stockholm, Sweden, April, 2010.

Iosifidis, Y.; Mallik, A.; Mamagkakis, S.; De Greef, E.; Bartzas, A.; Soudris, D. and Catthoor, F.: A framework for automatic parallelization, static and dynamic memory optimization in MPSoC platforms. In *Proceedings of the 47th ACM/IEEE Design Automation Conference – DAC*, Anaheim, USA (2010). Also listed in Software Synthesis, Code Generation deliverable and in Integration driven by Industrial Applications deliverable.

Baloukas, C.; Papadopoulos, L.; Soudris, D.; Stuijk, S.; Jovanovic, O.; Schmoll, F.; Cordes, D.; Pyka, R.; Mallik, A.; Mamagkakis, S.; Capman, F.; Collet, S.; Mitas, N. and Kritharidis, D. Mapping embedded applications on MPSoC - The MNEMEE approach. In *IEEE Computer Society Annual Symposium on VLSI – ISVLSI*, Lixouri, Greece (2010). Also listed in Software Synthesis, Code Generation deliverable and in Integration driven by Industrial Applications deliverable.

R. Santos, P. Pedreiras, M. Behnam, T. Nolte, L. Almeida. Hierarchical Server-Based Traffic Scheduling in Ethernet Switches. Demo at CRTS 2010 - 3rd Workshop on Compositional Theory and Technology for Real-Time Embedded Systems (satellite of RTSS 2010), San Diego, USA, November 29, 2010.

Rui Santos, Paulo Pedreiras, Farahnaz Yekeh, Thomas Nolte, Luis Almeida. On Hierarchical Server-based Communication with Switched Ethernet. (Work-in-Progress session) ETFA 2010, 15th IEEE Conference on Emerging Technologies and Factory Automation. Bilbao, Spain, 13-16 September 2010.

Rui Santos, Moris Behnam, Thomas Nolte, Luís Almeida and Paulo Pedreiras. Schedulability Analysis for Multi-level Hierarchical Server Composition in Ethernet Networks. RTN'10, 9th Workshop on Real-Time Networks, (satellite of ECRTS'10), Dublin, Ireland, July 2010.

Rui Santos, Paulo Pedreiras, Luis Almeida, Alexandre Vieira, Thomas Nolte, Ricardo Marau, Arnaldo Oliveira. Flexible, Efficient and Robust Real-Time Communication with Server-based Ethernet Switching. WFCS 2010 - 8th IEEE Workshop on Factory Communication Systems. Nancy, France. 19-21 May 2010.

R. Santos, A. Vieira, R. Marau, P. Pedreiras, A. Oliveira, L. Almeida, T. Nolte. Improving the efficiency of Ethernet switches for real-time communication. WARM 2010 - Workshop on Adaptive Resource Management (within the Cyber-Physical Systems Week), Stockholm, Sweden, 12 April 2010.

Vanessa Romero S., Karl-Erik Årzén, Stefan Schorr, Raphael Guerra, Gerhard Fohler, Johan Eker, Harald Gustafsson, Adaptive Resource Management for Mobile Terminals - The ACTORS approach, Proceedings of 1st Workshop on Adaptive Resource Management (WARM10), Stockholm, Sweden, April 2010.

-- The above are new references, not present in the Y2 deliverable --

2.5 Keynotes, Workshops, Tutorials

Keynotes

Professor L. Almeida. A Dynamic Scheduling Approach to Designing Flexible Safety-Critical Systems. Seminar at McMaster University, Hamilton, Canada, 4th June 2010.

Professor L. Almeida. Hierarchical Distributed Architectures for Autonomous Mobile Robots: A Case Study. Invited talk at the Singapore Polytechnic, 16 March 2010.

Professor Luis Miguel Pinho of the Polytechnic of Porto, Portugal gave the keynote talk "Real-time Programming Paradigms and Languages" at the 16th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA 2010).

Marisol García-Valls, ARTEMIS Summit (Held jointly with ITEA2), *Gennt, Belgium – October 26-27, 2010*.

Marisol García-Valls, Discussion Pannel on "Services vs components in real-time embedded systems modeling", International Workshop on Distributed Architecture modeling for Novel component based Embedded Systems, DANCE 2010 *Tozeur, Tunisia – May 30th, 2010*

Marisol García-Valls, WARM Workshop (Held jointly with RTAS 2010), *Stockholm, Sweden – April 12th, 2010*.

Workshops

DATE 2010 Friday Workshop, <http://conferenze.dei.polimi.it/depcp/2010/index.html>, Designing for Embedded Parallel Computing Platforms: Architectures, Design Tools, and Applications. Dresden, Germany, March 12, 2010.

Marisol García-Valls and Francisco Gómez Molinero Special Session on Reconfigurable Networked Embedded Systems, 15th International Conference on Emerging Technologies and Factory Automation, ETFA 2010. 13-16 September, Bilbao (SPAIN).

Alejandro Alonso-Muñoz and Marisol García-Valls, Special Session on Quality of Service and Resource Management, 15th International Conference on Emerging Technologies and Factory Automation, ETFA 2010. 13-16 September, Bilbao (SPAIN).

Rob Davis, Workshop on Real-Time Systems Open Problems, at ECRTS 2010.

-- The above is new material, not present in the Y2 deliverable --

3. Milestones, and Future Evolution

3.1 *Problems to be Tackled in Year 4* (Jan 2011 – Dec 2011)

The main cluster-wide problem to be tackled during the final 12 months of ArtistDesign is to continue to develop the wiki addressing the taxonomy of resource usage that was started during the second year. The objective of the wiki is to capture the policies, analysis techniques, and modelling approaches that are appropriate for resource management for a wide class of embedded systems. Both offline verification and online control are covered.

The other main focus is on multiprocessor and platforms and how they can be used in real industrial applications with timing constraints.

In addition, most of the 26 technical activities described in Section 2 will continue, as will the work addressing the milestones described below.

The work of this activity will continue to contribute to the Adaptivity and Predictability cross-cluster activities.

3.2 *Current and Future Milestones*

We continue to work on many of the milestones set during Years 1, 2 and 3:

- Extend the taxonomy of resource usage.
- Extend the use of hierarchical and contract-based scheduling to multi-resource systems.
- Produce effective scheduling and placement algorithms for multiprocessor systems.
- Produce mode change algorithms suitable for multiprocessor systems.
- Determine an effective way of undertaking (static) architectural tradeoffs.
- Determine an effective way of undertaking (dynamic) adaptive resource management, making use of feedback techniques from the control environment.
- Define a framework that can accommodate multiple time-frames within a single system and facilitate hierarchical scheduling, cascade control and other means of separating temporal concerns.

In year 4 some more specific topics have been identified:

- Dynamic priority assignment schemes for efficient, low-cost scheduling on multiprocessors.
- Building resource-awareness into communication middleware.
- Evaluation and possible enhancement of the expressive power of certain languages to integrate the support for QoS aware scheduling (e.g. real-time Java).
- Integration of support for bounded-time reconfiguration in networked embedded systems.
- Development of new theoretical results in the field of energy and power systems.

- Investigating the management of resources such as bandwidth and frequency in wireless networks.
- Resource management in distributed contexts.

3.3 Main Funding

The ArtistDesign NoE funds integration and building excellence with the partners, and with the European research landscape as a whole. Beyond this “glue” for integration and excellence, during Year 3 this activity has benefited from direct funding from:

ACTORS - EU STREP on adaptive resource management involving Ericsson, ULUND, TUKL, SSSA, EPFL, and Evidence.

LCCC - Lund Center for Control of Complex Engineering Systems is a Linneaus center at ULUND funded by the Swedish Research Council.

ELLIIT - The Linköping - Lund Strategic Research Center for IT and Mobile Computing funded by the Swedish funding agency VINNOVA.

CHAT - EU STREP on networked automation systems involving ULUND plus several non-ARTIST partners.

iLAND – mIddLewAre for deterministic dynamically reconfigurable Networked embedded systems – ARTEMIS funded project in which the following partners are involved: University Carlos III of Madrid (Technical Coordinator of iLAND project is Marisol García-Valls) and University of Porto.

MOSART IST-215244 Project: Mapping Optimization for Scalable multi-core ARchiTecture. ArtistDesign partners involved: Interuniversitair Micro-Elektronica Centrum (IMEC) vzw. and Kungliga Tekniska Hagskolan (KTH) <http://www.mosart-project.org/>

MNEMEE IST-216224 Project: Memory maNagEMEnt technology for adaptive and efficient design of Embedded systems. ArtistDesign partners involved: Interuniversitair Micro-Elektronica Centrum (IMEC) vzw. and Technische Universiteit Eindhoven (TU/e) <http://www.mnemee.org/>

2PARMA IST- 248716 Project: PARAllel PARadigms and Run-time MANagement techniques for Many-core Architectures. ArtistDesign partners involved: Interuniversitair Micro-Elektronica Centrum (IMEC) vzw. and RWTH Aachen University <http://www.2parma.eu/>

SCALOPES: is an ARTEMIS project on *SCalable LOW Power Embedded platformS*. ArtistDesign partners involved: Interuniversitair Micro-Elektronica Centrum (IMEC) vzw. and Universidad de Cantabria <http://www.scalopes.eu/>

REHEAT (Real-time scheduling on heterogeneous multicore architectures - PTDC/EIA-CCO/105716/2008), project funded by Portuguese Science Council (FCT).

RECOMP (Reduced Certification Costs for Trusted Multi-core Platforms - Artemis 100202), supported by the European Community through the Joint Undertaking ARTEMIS.

FlexWARE - Catania is involved in the IST 7FP STREP project “Flexible Wireless Automation in Real-Time Environments”, aiming at the implementation of a novel platform for the support of real-time communication over Wireless Local Area Networks based on the IEEE 802.11standard. Project webpage: www.flexware.at

MADES – EU funded project 1st April 2010 for 30 months involving TXT e-solutions (IT), Softeam (FR), University of York (UK), Politecnico di Milano (IT), The Open Group (UK), EADS Deutschland GmbH (DE). Aims to develop a UML / MARTE based model-driven approach for the design, validation, simulation, and code generation of complex embedded systems.

Tempo – ESPRC funded project in UK for York. Dealing with complex heterogeneous platforms including multicore and FPFAs.

GUARANTEE - A Guardian Angel for the Extended Home Environment. UPM is involved in this ITEA2 Project. The aim is to develop support for safety applications for reducing risks at home.
<http://www.guarantee-itea2.eu/>

RT-MODEL – Real-Time platforms for real-time embedded systems design based on models. UPM, UP Valencia and U Cantabria participate in this project. It is funded by the Spanish Education Ministry.

-- in the above list four projects started in 2010

4. Internal Reviewers for this Deliverable

- Robert Davis, University of York
- Michael Gonzalez Harbour, Universidad de Cantabria
- MARISOL GARCIA VALLS, Universidad Carlos III de Madrid