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

Activity Progress Report for Year 2

# 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 resources 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 | December 18th 2009 |

# Table of Contents

# 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.*

Dr Stylianos Mamagkakis (IMEC)

*The 'Runtime resource management for MPSoC' team at the Nomadic Embedded System 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.*

***Aveiro has been replaced by UnivPorto during the current year.***

## 1.2 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.*

Associate 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.*

Dr Liesbeth Steffens Steffens (NXP Semiconductors, Eindhoven – Affiliated to Kaiserslautern)
*The team is an industrial partner working with media processing for consumer electronics. Focus is on integrating multiple interdependent resources (processor, caches, memory bandwidth).*

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.*

***University of Aveiro, formerly a core partner, joined as affiliated participant during 2009.***

### 1.3 Starting Date, and Expected Ending Date

This activity started at the commencement of ArtistDesign and continues work been 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 1.***

### 1.4 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 1.***

## 1.5 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+ by 2010); 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 1.***

## 1.6 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 deallocation 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 significant changes with respect to Year 1.***

## *1.7 Work achieved in Year 1*

As indicated in the Description of Work the first milestone set for this activity during 2009 was to produce, in the second year, the outlines 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 Y1 deliverable, in sections 1.7 and 3.1.***

## *1.8 Problems Tackled in Year 2*

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: 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. 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 ARNIC 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 activity 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. This has not yet lead to a published paper, although 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 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 within York and is contributing to the taxonomy work. See TA1.

***The above is new material, not present in the Y1 deliverable.***

# 2. Summary of Activity Progress in Year 2

## *2.1 Technical Achievements*

**1. Development of the Taxonomy as a wiki (all partners lead by York)**
An initial definition of the taxonomy has been produced, but we were concerned that this resource should be useful to the wider community and be seen as a dynamic document that was continually modified and could become the means by which the key objectives of the activity could be focussed. To this end effort was placed to produce a wiki, to give it some initial structure and to start 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). It will also at some point 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). The variables indicate things like: how much does a specific task execute in this time interval. We can distinguish between the following types of constraints:

- Capacity Constraints - this battery stores energy of the amount 1kWh and it cannot be recharged; therefore in any time interval, the amount of energy drawn from this battery cannot exceed 1kWh;
- This flash-memory can be written to 100 000 times.
- Resource Management Constraints: on the use of a resource due to the inherent nature of the resource, for example
    - cumulative processing capacity cannot be used by a single task in parallel;
    - temperature of a CPU may not exceed 100 degrees;
    - limit of X number of rewrites of a block in flash;
- Dispatching policy, for example
    - the CAN bus, is shared in a non-preemptive fixed-priority fashion;
    - device driver for CAN queues transmission requests in FIFO order;
- Workload Constraints, for example,
    - Number of task arrivals in any given time interval
    - WCETs of tasks
    - Type of tasks: e.g. best effort vs (hard) real-time

Multiple resources consumption functions may have common influences. For example: Power consumption will be dependent on the frequency in a DVFS capable processor. However, changing the frequency does not only change the power consumption, but also the available execution cycles in a given release/deadline interval.

We note that time is not a resource in itself but it is used to “segment” a resource (such as energy in the battery) into sub-resources (energy supplied by the battery in a certain time interval). However, the number of processor cycles available in a time interval is a resource.

A resource (such as a processor) may also be segmented into other resources (such as register file, program counter, ALU, floating-point arithmetic units, L1 cache, branch-prediction logic, etc) but this is “space-segmentation” rather than “time-segmentation”.

**2. Hierarchical Scheduling (Porto, York, Cantabria, UnivPorto, Aveiro, Malardalen)**

Within the activity a number of tasks were undertaken to tackle issues arising in this context. One approach deals with reducing the interaction of system integrators with subsystem vendors in scheduling resource-sharing applications in an ARINC 653-2 environment. This is achieved by modelling the ARINC framework as a two-level hierarchical scheduler and providing appropriate interface generation tools to the subsystem vendors. In the other work, we focus on theoretical foundations for schedulability analysis of hierarchical systems. We define notions of optimality for subsystem interfaces with respect to resource usage and also classify existing interface generation techniques based on these notions.

This work on hierarchical scheduling also forms a key part of the taxonomy.

Cantabria has contributed by extending the MAST model used to describe the timing behaviour of real-time systems with partitioned hierarchical scheduling facilities such as those defined in the ARINC 653 standard. MAST is a modelling framework that can be used to model real-time applications and apply automatic schedulability analysis tools. A new version of the model is being produced, called MAST-2, which will align the names of the modelling elements to the new MARTE UML profile for real-time embedded systems, and will add support for partitioned systems, both in the processors and in the networks, as well as support for deterministic Ethernet switches according to the AFDX industry standard.

UnivPorto and Aveiro continued evolving the FTT-SE framework that allows creating deterministic partitions over general purpose Ethernet switches that can be adapted and reconfigured on-line. Together with Malardalen, this framework was extended with a new protocol, Server-SE, that allows providing transparent partitions that can be used by any legacy application without need for adaptation, using any available CPU scheduling policy. Finally, this framework has recently been enriched with the integration of the FTT master inside the switches, which are now capable of managing the network resources in the time domain, independently from the end nodes boosting the system robustness and integration capabilities.

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

Standardisation work continues for 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. A particular focus this year has been the support that languages can provide for multiprocessor platforms, in particular multicore schemes. To this end an ArtistDesign supported workshop was held (in September in Italy) on real-time language issues. From this workshop a number of possible language modifications were developed and a set of these have been forwarded to the Ada standard’s working group on which ArtistDesign members participate. The findings have also been fed into the Real-Time Java standardisation community on which an ArtistDesign member also participates. It is hoped that during the coming year these proposals will be adopted.

**4. Memory Resource Management (Universidad Politécnica de Valencia, York, NXP, Kaiserslautern).**

Dynamic memory in real-time systems has been avoided due to two fundamental problems: allocation and deallocation in bounded time, and the fragmentation problem. Recent research

results have removed the unbounded timing behaviour of the dynamic memory allocation. TLSF is a fast and constant time memory allocator. This work addresses the issues of handling dynamic memory in real-time systems. Dynamic memory management has a strong impact in embedded systems with resource constraints allowing for an increase in the flexibility and adaptability of the applications. This work is focused on providing: i) a vision of the memory as a resource in the same way that other resources are considered; ii) a memory model and a memory reservation architecture being able to manage the spare memory of the system and, iii) an acceptance test for memory contracts and a memory reclaiming mechanism of the memory associated requests.

NXP and TUKL continued work started during ARTIST2 on integrating real-time scheduling and cache management on multiprocessor platforms. A joint PhD student is supported by this activity, but has stopped PhD work due to restructuring within NXP.

**5. Deadline-Period analysis in embedded control systems (Universidad Politécnica de Valencia)**

In the design of a real-time application it is fundamental to know how a change in the task parameters would affect the feasibility of the system. Relaxing the classical assumptions on static task sets with fixed periods and deadlines can give higher resource utilisation and better performance. But the changes on task parameters have to be done always maintaining feasibility. In practice, period and deadline modifications are only necessary on single tasks. Our work focuses on finding the feasibility region of deadlines and periods (called D-P feasibility region) for a single task in the context of dynamic, uniprocessor scheduling of hard real-time systems. This way, designers can choose the optimal deadline and period pairs that best fit application requirements. We provide an exact and an approximated algorithm to calculate this region.

**6. Resource management for control tasks (UPC, UCSC)**

In many application areas, including control systems, careful management of system resources is the key to providing the best application performance. Traditional control systems with multiple control loops statically allocate a fixed portion of the system resources to each controller based on their average or worst-case resource requirements. However, controllers' resource needs vary depending on the jobs they perform and the state of the systems they control. A controller of a plant operating close to its equilibrium requires fewer resources than a controller of a plant operating far from its equilibrium point. The Draco dynamic rate control system exploits this fact by dynamically allocating resources to control systems based on system state. The research demonstrates that Draco provides significantly better overall control performance with much less resources than static controllers.

Event-driven control systems provide interesting benefits such as reducing resource utilization. A novel formulation for the design of event-driven controllers has been developed to allow a trade-off between their resource demands and their control performance.

**7. Extending resource reservation to multicore systems (Pisa)**

Theoretical work has been done by the Scuola Superiore Sant'Anna of Pisa to extend the resource reservation programming paradigm to multicore systems. In particular, a new scheduling framework has been investigated to model, analyze, and partition parallel real-time applications on a multicore platform. The method handles sporadic real-time applications with timing and precedence constraints and allocates them to a set of virtual processors to make the allocation process independent of the physical platform. Then each virtual processor is

implemented using a server mechanism that is mapped to a physical core that can guarantee the reserved bandwidth.

**8. Limited preemptive scheduling (Pisa)**

Research has been undertaken by the Scuola Superiore Sant'Anna of Pisa on limited preemptive scheduling, to investigate the possibility of reducing cache-related preemption delays in task worst-case execution times (WCETs), so reducing variability in WCET estimations and increasing the predictability of real-time system without loosing efficiency in resource usage.

**9. Produce effective scheduling and placement algorithms for multiprocessor systems (Porto)**

A multicore processor is interesting from a resource-management perspective is that a multicore processor has resources that are shared between processor cores. This includes (i) a shared interconnection network between the processor core and the main-memory modules and (ii) potentially a cache memory that is located inside the multicore chip but is shared between the processor cores.

The Polytechnic institute of Porto has been working on improving the utilisation bounds of multicores by reducing the number of preemptions using the concept of notional processors. Further work considers the clustering of cores on multichip multicores, to avoid the more expensive off chip migration of tasks. The issue of resource sharing and mutual exclusion on a macro level (i.e. explicit sharing of resources like coprocessors, but not transparent sharing like memory buses) has been addressed in our work on scheduling under mutual exclusion constraints. We investigated techniques for clustering of tasks using a novel concept of virtual processors. The physical multicore platform is virtualised using a global scheduling strategy, and virtual resources are then allocated to task clusters using a local cluster-specific scheduling strategy.

Finally, a generalisation of the priority ceiling protocol for global fixed priority scheduling on multiprocessors has been investigated.

**10. Effective way of undertaking (dynamic) adaptive resource management (Porto, UnivPorto, Aveiro, Valencia, UC3M)**

The growing complexity and dynamism of many embedded application domains (including consumer electronics, robotics, automotive and telecommunications), makes it increasingly difficult to react to load variations and adapt the system's performance in a controlled fashion within an useful and bounded time. This is particularly noticeable when intending to benefit from the full potential of an open distributed cooperating environment, where service characteristics are not known beforehand and tasks may exhibit unrestricted QoS inter-dependencies. In this context, the Polytechnic Institute of Porto has been addressing cooperative service execution, allowing resource constrained devices to collectively execute services with more powerful neighbours. Nodes dynamically group themselves into coalitions, allocating resources to each new service and establishing an initial service configuration which considers the satisfaction of the QoS constraints associated with the new service and the impact on the global QoS caused by the new service's arrival.

The dynamic changing of the set of services being executed also impacts in the serviced configuration. Therefore, in this year, the research focused on both centralised and decentralised coordination models for distributed QoS adaptation based on a simple and effective feedback mechanism to reduce the complexity of the needed interactions among

inter-dependent nodes of a cooperative distributed system. The latter was followed by Porto while the former was taken by UnivPorto, Aveiro and Valencia, based on the FTT-SE protocol, and it was applied to a distributed mutlimedia setup with dynamic number of streams and dynamic scenarios, using feedback control to manage the compression of each stream and a greedy approach to distribute or request bandwidth.

Another angle at open systems and required reconfiguration of system parameters has been taken with a temporal isolation approach integrating multiple resource constraints. This work not only deals with the dynamics of an open system but in particular with dynamic slack management and the requirements of underlying hardware like cache-related preemption delay often ignored by existing approaches. Special consideration is taken for the integration and interdependencies of approaches dealing with problems in isolation.

### 11. Mode change algorithms suitable for multiprocessor systems (Porto)

The problem of global and preemptive scheduling of multi-mode real-time systems upon identical multiprocessor platforms has been tackled by the Polytechnic Institute of Porto. Since the system considered is a multi-mode system, the system can change from one mode to another such that the current task set is replaced with a new task set. Ensuring that deadlines are met requires not only that a schedulability test is performed on tasks in each mode but also that (i) a protocol for transitioning from one mode to another is specified and (ii) a schedulability test for each transition is performed. This work proposes two protocols which ensure that all the expected requirements are met during every transition between every pair of operating modes of the system. Moreover, the correctness of the proposed algorithms is proven by extending the theory about the makespan determination problem.

### 12. Virtual Educational Laboratory of Distributed Process Control (Catania)

Education in the field of resource management in distributed process control and intelligent actuation and measurement (IAM) systems can benefit from the adoption of tools allowing students to test and compare several system configurations and to assess the system behaviour under different settings and operating conditions. The team working at the University of Catania tackled this problem by realizing an educational virtual laboratory that permits interaction with the simulator of a control process executed on a remote server. The simulator is a software copy of the physical process that was realized in the IAM-Pilot, a research pilot plant realized at the ENEL (the former Italian National Electric Energy Company) Research Centre, Milan, with the aim of testing concepts regarding distributed systems, technologies, and components. Such a simulator implements a closed circuit of treated water implementing four regulation loops for pressure, level, flow, and temperature, respectively, and offers a variety of control processes and field components. The tool produces results that allow users to achieve a better understanding of the design problems typical of such systems. Users can configure the characteristics of hydraulic circuit and the devices used, and this allows them to practically verify the results that they obtained analytically. Moreover, the tool also permits them to evaluate the impact of some events, such as network overloading or deterioration of a valve, which are otherwise difficult to calculate. By monitoring both the behaviour of the system and the performance offered by the control system it is possible to evaluate the correctness of the work done.

### 13. Optimal period selection and scheduling for embedded controllers (ULUND, Linköping Univ)

The work on how to optimally schedule a set of distributed control algorithms studied by ULUND and Linköping University and introduced in the Y1 deliverable has continued 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. The Jitterbug toolbox from Lund is used to evaluate the control performance, taking the delay and jitter into account.

### 14. Sporadic Event-Based Control (ULUND, SSSA)

The work on event-based control performed mainly by ULUND continues. This year the work has focused on applications. Traditional linear time-invariant (LTI) control design assumes that measurements are taken at regular time intervals and have independent additive noise. A common practical case that violates this assumption is the use of encoders that give quantized position measurements; when the quantization is appreciable the measurement noise is far from LTI. A simple event-based controller based on simplifying a joint maximum *a posterior* estimator has been developed. This has been applied to a moving cart with quantized position measurements. The payoff for implementing the somewhat more complex event-based controller is to drastically reduce the effect of quantization noise in the experiments. In collaboration with SSSA new schemes for event-based PID control have also been developed.

### 15. Task concurrency management and access scheduling for dynamic systems (IMEC vzw., Univ. Complutense Madrid-UCM, National Technical Univ. of Athens-NTUA, Norwegian Univ. of Science and Technology-NTNU)

In the context of this collaboration, software metadata were extracted from embedded software applications mapped on MPSoC platforms via profiling and simulation. This metadata information was used in order to dynamically reconfigure the DMA access programming and the OS scheduler for lower power design and higher performance. Additionally, trade-offs between the optimizations was explored.

UC3M also works on achieving effective QoS-based resource management in mobile operating systems taking advantages of their concurrent task execution model.

### 16. Techniques to avoid beacon collisions in IEEE 802.15.4 Cluster-Tree Networks (Catania)

In cluster-tree topologies, IEEE 802.15.4 networks comprise multiple coordinators. When the beacon-enabled MAC mode is used, all the coordinators periodically generate beacon frames to synchronize the nodes belonging to their cluster. As a result, if the transmission of the beacon frames is not properly synchronized, a beacon frame may collide either with other beacon frames from different coordinators or with data frames from different clusters. Nodes not receiving beacon frames may lose the synchronization with their coordinator and thus get disconnected from the network. To solve this problem, the 2006 revision of the IEEE 802.15.4 standard introduced the support for the time-division superframe scheduling, which schedules the superframes by adding a time offset between the superframes of the parent and child coordinator. While this technique is able to avoid beacon collisions, it limits the network scalability, as no parallel communication is allowed unless coordinators are distant enough not to collide. The team at the University of Catania addressed the problem by introducing a novel technique to schedule the superframes of cluster-tree IEEE 802.15.4 networks over multiple channels. Such technique has been exploited by a novel scheduling algorithm, called a Multichannel Superframe Scheduling (MSS) that allows multiple clusters to schedule their superframes simultaneously on different radio channels. In this way, it is possible to schedule

sets of superframes which could not be schedulable under single-channel superframe scheduling algorithms.

**17. Multiresource contract-based scheduling framework (Cantabria, York, Malardalen, Pisa, Valencia, and other academic and industrial partners)**

The FRESCOR EU project was successfully completed during the reporting period achieving all of its main initial objectives. The FRESCOR project has produced a resource reservation framework that, as such, provides protection among the different software components running on top of it, facilitating the independence of their development and execution. But the framework produced in FRESCOR goes well beyond the capabilities of other resource reservation frameworks by providing adaptive reservations that can make use of spare capacity available; management of QoS parameters expressed at the application level; an integrated management of multiple schedulable resources including CPUs, networks, memory, disk bandwidth, and energy; management of time-protected shared objects; integration with component-based design methods; off-line schedulability analysis and simulation tools that allow the application developer to reason about the timing behaviour of the application before it is built; and an API that makes the application independent of the underlying operating system, networks and scheduling policies.

With the results achieved we can conclude that the adaptive resource reservation framework developed in the FRESCOR project is an enabling technology that can contribute to managing the complexity of developing real-time embedded applications. The contribution is on the management of resources to meet timing requirements and its usage will contribute to the ability to develop more complex systems capable of achieving more functionality, the reliability of these systems by providing a timely management of the resources, and the reduction of the cost of development by integrating the technology in the design flow and by raising the level of abstraction of the management of resources.

**18. Dataflow Scheduling using Constraint Programming (ULUND, Ericsson)**

Within the ACTORS project ULUND and Ericsson are applying constraint programming for static uni-processor scheduling of synchronous data-flow (SDF) networks and cyclo-static data flow (CSDF) networks. With the current focus on multi and many-core platforms data-flow modeling has become increasingly popular because of its explicit support for parallelism. The CAL programming language is an actor language especially aimed at media streaming applications. Using CAL it is possible to obtain networks that are SDF, CSDF, as well as dynamic. For performance reasons it is an advantage if one can automatically detect sub-networks that are SDF or CSDF, since in that case it is possible to statically schedule these parts at design-time. In the current work, automatic methods are derived for classifying actors and finding statically schedulable sub-networks. Once these are available constraint programming is used to calculate repetition vectors and static schedules that, e.g., minimize the internal buffer requirements. Once these are available, the involved actors are automatically merged into composite actors according to the generated schedule. The constraint programming language JaCoP is used together with the miniZinc language for expressing finite domain optimization problems.

**18. Low-level adaptation mechanisms in contract-based systems (UC3M, UPM)**

Consumer electronics platforms are progressively being introduced in most domains, from personal appliances to high-quality multimedia control infrastructures. Multimedia applications compete for resources in a greedy manner, which may threaten the safe system execution. Arbitration of application execution is needed to keep system stability, especially in the event of

unexpected changes in the resource requirements demanded by applications. Carlos III University of Madrid (UC3M) and Technical University of Madrid (UPM) are collaborating to define mechanisms for ensuring that changes in the system are handled in a smooth way and keeping the output QoS stable.

This work has resulted in mode change algorithms for applications that do not tolerate data losses during their execution, the integration of safety bands for filtering and handling peaks in resource demands, and resource application characterization and a contract model that facilitates a predictable behavour, even during reconfiguation.

***The above is new material, not present in the Y1 deliverable.***

## *2.2 Individual Publications Resulting from these Achievements*

**University of York**

F. Zhang and A. Burns**,** Schedulability Analysis for Real-Time Systems with EDF Scheduling**,** IEEE Transactions on Computers, 58(9), 1250-1258, 2009.

R. I. Davis, T. Rothvo, S. K. Baruah and A. Burns, Exact Quantification of the Sub-optimality of Uniprocessor Fixed Priority Pre-emptive Scheduling, Real Time Systems Journal, 43(3), 211-258, 2009.

A. Zuhily and A. Burns, Exact Scheduling Analysis of Non-Accumulatively Monotonic Multiframe Tasks, Real-Time Systems Journal, 43, 119-146, 2009.

R. Davis and A. Burns, Robust priority assignment for messages on Controller Area Network (CAN), Real-Time Systems, 41(2), 152-180, 2009.

K. Yu and N. Audsley, A Mixed Timing System-Level Embedded Software Modelling and Simulation Approach, Proc. of International Conference on Embedded Software and Systems, ICESS, 2009.

M. Alrahmawy and A.J. Wellings, An RTSJ-based reconfigurable server component, Proceedings of the 7th International Workshop on Java Technologies for Real-Time and Embedded Systems, 31-40, 2009.

M-S. Kim and A.J. Wellings, Applying fixed-priority preemptive scheduling with preemption thresholds to asynchronous event handling in the RTSJ, Proceedings of the 7th International Workshop on Java Technologies for Real-Time and Embedded Systems, 81-89, 2009.

A. Burns, A.J. Wellings and F. Zhang, Combining EDF and FP Scheduling: Analysis and Implementation in Ada 2005, Proceedings of Reliable Software Technologies - Ada-Europe 2009, LNCS 5570, 119-133, 2009.

F. Zhang and A. Burns, Improvement to Quick Processor-demand Analysis for EDF-Scheduled Real-Time Systems, Proceedings of the 21st Euromicro Conference on Real-Time Systems (ECRTS), 76-86, 2009.

T. Richardson and A.J. Wellings, Providing Temporal Isolation in the OSGi Framework,7th International Workshop on Java Technologies for Real-Time and Embedded Systems, 2009.

R. Davis and A. Burns, Quantifying the Sub-optimality of Uniprocessor Fixed Priority Pre-emptive Scheduling for Sporadic Tasksets with Arbitrary Deadlines, Real-Time Networks and Systems Conference, Paris, 2009.

Z. Shi and A. Burns, Real-Time Communication Analysis with a Priority Share Policy in On-Chip Networks, 21st Euromicro Conference on Real-Time Systems (ECRTS), 1-10, 2009.

Z. Shi and A. Burns, Improvement of Schedulability Analysis with a Priority Share Policy in On-Chip Networks, 17th International Conference on Real-Time and Network Systems (RTNS), Paris, 75-84, 2009.

F. Zhang, A. Burns and S. Baruah, Sensitivity Analysis for Real-Time Systems, University of York, Computer Science Dept, YCS-2009-438, 2009.

**Scuola Superiore Sant'Anna, Pisa**

Enrico Bini, Giorgio Buttazzo and Yifan Wu, "Selecting the Minimum Consumed Bandwidth of an EDF Task Set", Proceedings of the 2nd Workshop on Compositional Theory and Technology for Real-Time Embedded Systems (CRTS 2009), Washington, D.C., USA, December 1, 2009.

Enrico Bini and Giorgio Buttazzo, "The space of EDF deadlines: the exact region and a convex approximation", Real-Time Systems, Vol. 41, No. 1, pp. 27-51, January 2009.

Enrico Bini, Giorgio Buttazzo and Marko Bertogna, "The Multi Supply Function Abstraction for Multiprocessors", Proceedings of the 15th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA 2009), Beijing, China, August 24-26, 2009 (Best Paper).

Gang Yao, Giorgio Buttazzo and Marko Bertogna, "Bounding the Maximum Length of Non-Preemptive Regions Under Fixed Priority Scheduling", Proceedings of the 15th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA 2009), Beijing, China, August 24-26, 2009.

Enrico Bini, Giorgio Buttazzo, and Giuseppe Lipari, "Minimizing CPU energy in real-time systems with discrete speed management", ACM Transactions on Embedded Computing Systems, Vol. 8, Issue 4, July 2009.

Marko Bertogna, Michele Cirinei, Giuseppe Lipari. "Schedulability analysis of global scheduling algorithms on multiprocessor platforms", IEEE Transactions on Parallel and Distributed Systems. 20(4): 553-566. April 2009.

Marko Bertogna. "Evaluation of existing schedulability tests for global EDF", Proceedings of Real-time systems on multicore platforms: Theory and Practice, (in conjunction with ICPP 2009), Vienna, Austria, September 2009.

Yifan Wu and Marko Bertogna. "Improving Task Responsiveness with Limited Preemptions", Proceedings of 14th IEEE International Conference on Emerging Techonologies and Factory Automation, Palma de Mallorca, Spain, September 2009.

Dario Faggioli, Michael Trimarchi, Fabio Checconi, Marko Bertogna, Antonio Mancina. "An Implementation of the Earliest Deadline First Algorithm in Linux", Proceedings of the 24th Annual ACM Symposium on Applied Computing, Honolulu, Hawaii, USA. March 2009.

**Universidad de Cantabria**

Héctor Pérez Tijero, J. Javier Gutiérrez and Michael González Harbour. “Support for a real-time transactional model in distributed Ada”. 14th International Real-Time Ada Workshop (IRTAW-14), October 7-9 in Portovenere, Italy

Mario Aldea Rivas and Michael González Harbour. “Execution time monitoring and interrupt handlers”. 14th International Real-Time Ada Workshop (IRTAW-14), October 7-9 in Portovenere, Italy

H. Pérez Tijero and J. J. Gutiérrez. Experience in Integrating Interchangeable Scheduling Policies into a Distribution Middleware for Ada. SIGAda 2009 ACM Annual International Conference on Ada and Related Technologies: Engineering Safe, Secure, and Reliable Software. St. Petersburg, Florida, USA, November 1-5, 2009 (awarded Best Paper).

Mario Aldea Rivas, Michael González Harbour and José F. Ruiz. “Implementation of the Ada 2005 Task Dispatching Model in MaRTE OS and GNAT”. Proceedings of the 14th International Conference on Reliable Software, LNCS Volume 5570/2009, june, 2009, ISBN:0302-9743, pp. 105,118.

**TUKL**

Raphael Guerra, Gerhard Fohler, A Gravitational Task Model with Arbitrary Anchor Points for Target Sensitive Real-Time Applications, Real-Time Systems Journal, September 2009.

Raphael Guerra, Gerhard Fohler, On-line Scheduling Algorithm for the Gravitational Task Model, ECRTS09 - 21th Euromicro Conference on Real-Time Systems, IEEE Computer Society, Dublin, Ireland, July 2009.

**Lund University**

T. Henningsson and A. Cervin, Comparison of LTI and Event-Based Control for a Moving Cart with Quantized Position Measurements, Proceedings of the European Control Conference, Budapest, Hungary, 2009.

**University of Aveiro / University of Porto**

Arnaldo Oliveira, Luis Almeida, Antonio B. Ferrari. The ARPA-MT Embedded SMT Processor and its RTOS Hardware Accelerator. IEEE Transactions on Industrial Electronics (to appear).

R. Santos, R. Marau, Alexandre Vieira, Paulo Pedreiras, Arnaldo Oliveira, Luis Almeida. A Synthesizable Ethernet Switch with Enhanced Real-Time Features. IECON 2009, 35th Annual Conf of the IEEE Industrial Electronics Society, Porto, Portugal, 3-5 Nov 2009.

**Polytechnic Institute of Porto**

A Compositional Scheduling Framework for Digital Avionics Systems, Arvind Easwaran, Insup Lee, Oleg Sokolsky, Steve Vestal in 15th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA 2009), Beijing, China.

Simulation Relations, Interface Complexity, and Resource Optimality for Real-Time Hierarchical Systems, Insup Lee, Arvind Easwaran, Madhukar Anand, Anna Philippou and Oleg Sokolsky in Proceedings of Workshop on Reconciling Performance with Predictability 2009, October 15, 2009, Grenoble, France.

Notional processors: an approach for multiprocessor, Konstantinos Bletsas, Björn Andersson, Proceedings of the 15th IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS'09), 2009.

Optimal virtual cluster-based multiprocessor scheduling, Arvind Easwaran, Insik Shin, Insup Lee, Real-Time Systems Journal, ISSN 0922-6443, Volume 43, Number 1, September, 2009.

Scheduling Sporadic Tasks on Multiprocessors with Mutual Exclusion Constraints, Arvind Easwaran, Björn Andersson, Workshop on Real-time Systems on Multicore Platforms: Theory and Practice (held in conjunction with ICPP 2009), Vienna, Austria, 2009.

Preemption-light multiprocessor scheduling of sporadic tasks with high utilisation bound, Konstantinos Bletsas, Björn Andersson, in Proceeding of the 30th IEEE Real-Time Systems Symposium (RTSS 2009), 1-4 December 2009, Washington (DC), USA.

Resource Sharing in Global Fixed-Priority Preemptive Multiprocessor Scheduling, Arvind Easwaran and Bjorn Andersson, in Proceeding of the 30th IEEE Real-Time Systems Symposium (RTSS 2009), 1-4 December 2009, Washington (DC), USA.

Two Protocols Without Periodicity for the Global and Preemptive Scheduling Problem of Multi-Mode Real-Time Systems upon Multiprocessor Platforms, Vincent Nelis, Joel Goossens, Björn Andersson, in 21st Euromicro Conference on Real-Time Systems (ECRTS 09), University of Dublin, Trinity College, July 1-3, 2009, Dublin, Ireland.

Nogueira, L., Pinho, L., "Time-bounded Distributed QoS-Aware Service,Configuration in Heterogeneous Cooperative Environments", Published in Journal of Parallel and Distributed Computing 69, pp. 491-507, 2009.

Nogueira, L., Pinho, L., Coelho, J., "Coordinated Runtime Adaptations, in Cooperative Open Real-Time Systems", Published in the 7th IEEE/IFIP International Conference on Embedded and Ubiquitous Computing, 2009.

Towards Real Multi-Criticality Scheduling, Stefan M. Petters, Martin Lawitzky, Ryan Heffernan, Kevin Elphinstone, in 15th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA 2009), Beijing, China, 2009.

Björn Andersson, A Pseudo-Medium-Wide 8-Competitive Interface for Two-Level Compositional Real-Time Scheduling of Constrained-Deadline Sporadic Tasks on a Uniprocessor, in Proceedings of the 2nd Workshop on Compositional Theory and Technology for Real-Time Embedded Systems (held in conjunction with RTSS), December 1, 2009, Washington DC, USA.

**IMEC**

Couvreur, C.; Nollet, V.; Catthoor, F. and Corporaal, H.: Fast multi-dimension multi-choice Knapsack heuristic for MP-SoC run-time manangement. In Journal ACM Transactions on Embedded Computing Systems (to appear in 2009)

Obermaisser, R.; Kopetz, H.; Huber, B.; Goedecke, M.; Warris, H.; Couvreur, C.; Gherman, V.; Ovaska, E.; Campos, S. and Paukovits, C.: GENESYS - A candidate for an ARTEMIS cross-domain reference architecture for embedded systems. In book chapter to appear through publisher SVH under ISBN 978-3-8381-1040-0, p.202, Saarbrücken, Germany, 2009.

**Universidad Politécnica de Valencia**

P. Balbastre, I. Ripoll, A. Crespo, Period sensitivity analysis and D-P domain feasibility region in dynamic priority systems. Journal of Systems and Software 82(7): 1098-1111, 2009.

**University Carlos III of Madrid**

M. García-Valls, P. Basanta-Val, and I. Estévez Ayres. Concurrency Programming Models in Real-Time Mobile Platforms. In Proc. of the 4th International IEEE Service Oriented Architectures in Converging Networked Environments. (IEEE SOCNE'09) Bradford, UK. May 2009.

I. Estévez Ayres, P. Basanta-Val, M. García-Valls, J. Arias-Fisteus and Luis Almeida. QoS-aware Real-Time Composition Algorithms for Service-Based Applications. August 2009. IEEE Trans. on Industrial Informatics.

**University of Catania**

E. Toscano, L. Lo Bello, “A Multichannel Approach to Avoid Beacon Collisions in IEEE 802.15.4 Cluster-Tree Industrial Networks”, In Proceedings of the 14th IEEE International Conference on Emerging Technologies and Factory Automation, ETFA'09, Sept.22-26, Palma de Mallorca, Spain, 2009.

L. Lo Bello, O. Mirabella, A. Raucea, L. Capetta, “ENEL PILOT: From a Research Testbed to a Virtual Educational Laboratory”, IEEE Transactions on Industrial Electronics, Vol.56, 2009.

**University of Pavia**

Marco Luigi Della Vedova, Tullio Facchinetti, Antonella Ferrara, and Alessandro Martinelli, "Real-time platooning of mobile robots: design and implementation", Proceedings of the 14th IEEE International Conference on Emerging Techonologies and Factory Automation (ETFA), Palma De Mallorca, Spain, 22-26, September, 2009.

Marco Luigi Della Vedova, Tullio Facchinetti, Antonella Ferrara, and Alessandro Martinelli, "Visual interaction for real-time navigation of autonomous mobile robots", Proceedings of the International Conference on CYBERWORLDS, University of Bradford, UK, September, 2009.

Roberto De Lotto, Tullio Facchinetti, Paolo Gamba, and Emanuele Goldoni, "Wireless Sensor Networks for planning processes: applications and case study", in Planning, Complexity and New ICT, Rabino and Caglioni editors, Alinea editions, pp. 127-136, September, 2009.

Roberto De Lotto, Tullio Facchinetti, Paolo Gamba, and Emanuele Goldoni, "Wireless Sensor Networks for monitoring urban environments: evaluation and practical considerations", in Planning, Complexity and New ICT, Rabino and Caglioni editors, Alinea editions, pp. 137-146, September, 2009.

***The above are new references, not present in the Y1 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 symposium RTSS and the European conference on real-time issues, ECRTS. Interactions also take place as part of other funded

EU projects such as FRESCOR which completed during 2009. Finally ArtistDesign specific meetings have taken place, for example a meeting in Pisa. At this meeting the framework for the taxonomy on resource usage was progressed and the decision to set up a wiki was made. Other meetings have been between smaller groups of partners. A number of PhD students have visited other labs both within ArtistDesign and beyond the project, for example Raphael Guerra from Kaiserslautern, visiting University of North Carolina, Chapel Hill, USA Oct - Dec 2009 to work on multiprocessor scheduling issues.

The Computer Engineering and Networks Laboratory at ETH Zurich, Switzerland, the Institute of Computer and Communication Network Engineering at TU Braunschweig, Germany, and the University of Cantabria have collaborated in a study of different real-time analysis methods for distributed systems evaluating their influence on the results of the analysis.

UnivPorto, Aveiro and Malardalen continue developing efforts in network partitioning management, particularly concerning the use of CPU server-based scheduling techniques in Ethernet switches.

ULUND, SSSA, TUKL, EPFL, Ericsson, and Evidence continues their collaboration on feedback-based resource scheduling for multimedia terminals within the EU STREP project ACTORS (Adaptivity and Control of Resources in Embedded Systems). Similar efforts in the scope of industrial monitoring systems continue being pursued by UnivPorto, Aveiro and Valencia.

ULUND and Linkoping University collaborate on integrated scheduling and synthesis of distributed and embedded control applications.

UC3M and UPM have continued with their collaboration on mechanisms for low-level mechanisms in contract-based systems.

NXP collaborates with a number of internal partners and external groups such as: ST Microelectronics, Fraunhofer Heinrich-Hertz-Institut, INCORAS, Politecnico di Milano, INRIA in the EU STREP project Open Media Platform on implementing memory and CPU resource management techniques for the OpenMax streaming framework API.

Overall many more joint publications have been produced in year 2 compared with the previous year.

***No significant changes to how the activity has progressed its work during Year 2. Significantly more joint publications have been produced.***

## *2.4 Joint Publications Resulting from these Achievements*

I. Ripoll, P. Balbastre, M. Masmano, A. Crespo and A. Burns, Contract based management of the memory resource, 17th International Conference on Real-Time and Network Systems (RTNS 2009), October 26-27, Paris, 2009.

A. Zabos, R. I. Davis, A. Burns and M. González Harbour, Spare Capacity Distribution Using Exact Response-Time Analysis, 17th International Conference on Real-Time and Network Systems (RTNS), 97-106, October 26-27, Paris, 2009.

P. Martí, M. Velasco and E. Bini. The Optimal Boundary and Regulator Design Problem for Event-Driven Controllers. In 12th International Conference on Hybrid Systems: Computation and Control (HSCC09), San Francisco, CA, USA, April 2009.

P. Martí, C. Lin, S. Brandt, M. Velasco and J.M. Fuertes. Draco: Efficient Resource Management for Resource-Constrained Control Tasks. In IEEE Transactions on Computers, Vol. 58, N. 1, Jan. 2009.

Kai Huang, Luca Santinelli, Jian-Jia Chen, Lothar Thiele, and Giorgio Buttazzo, "Adaptive Dynamic Power Management for Hard Real-Time Systems," Proceedings of the 30th IEEE Real-Time Systems Symposium (RTSS 2009), Washington D.C., December 2-4, 2009.

Kai Huang, Luca Santinelli, Jian-Jia Chen, Lothar Thiele, and Giorgio Buttazzo, "Periodic Power Management Schemes for Real-Time Event Streams," Proceedings of the 48th IEEE Conference on Decision and Control (CDC 2009), Shanghai, China, December 16-18, 2009.

Marko Bertogna, Nathan Fisher, Sanjoy Baruah. "Resource-sharing servers for Open Environments", IEEE Transactions on Industrial Informatics. 5(3): 202-220. August 2009.

Marko Bertogna, Nathan Fisher, Sanjoy Baruah. "Resource holding times: Computation and Optimization", Real-Time Systems: The International Journal of Time-Critical Computing, 41(2): 87-117. February 2009.

Thi Huyen Châu Nguyen, Pascal Richard, Enrico Bini, “Approximation techniques for response-time analysis of static-priority tasks”, Real-Time Systems, 43(2), pp. 147-176, October 2009.

Enrico Bini, Thi Huyen Châu Nguyen, Pascal Richard, Sanjoy Baruah, “A Response Time Bound in Fixed-Priority Scheduling with Arbitrary Deadlines”, IEEE Transactions on Computers 58 (2), pp. 279-286, February 2009.

Enrico Bini, Claudio Scordino, “Optimal Two-Level Speed Assignment for Real-Time Systems”, International Journal of Embedded Systems (IJES), Vol.4, No. 2, Pages 101-111, 2009.

Di Natale, Marco, Wang, Guoqiang and Sangiovanni-Vincentelli, Alberto, “Improving the size of communication buffers in synchronous models with time constraints”, IEEE Transactions on Industrial Informatics, Special Section on Real-Time Networked Embedded Systems, Volume 5, Number 3, August 2009

Sangiovanni-Vicentelli, A.; Di Natale, M., “Challenges and Solutions in the Development of Automotive Systems”, IEEE Transactions on CAD, Special Issue on Automotive systems, July 2009.

A. Ferrari, M. Di Natale, G. Gentile and P. Gai, “Time and Memory Tradeoffs in the Implementation of AUTOSAR Components”, DATE Conference, April 20-24 2009, Acropolis, Nice, France

Guoqiang Wang, Marco Di Natale, Pieter J. Mosterman, Alberto Sangiovanni-Vincentelli, “Automatic Code Generation for Synchronous Reactive Communication, International Conference on Embedded Software and Systems (ICESS) 2009, Hangzhou, China, 25-27 May 2009.

Qi Zhu, Yang Yang, Eelco Scholte, Marco Di Natale and Alberto Sangiovanni-Vincentelli, “Optimizing Extensibility in Hard Real-time Distributed Systems”, 15th IEEE Real Time and Embedded Technology and Applications Symposium (RTAS'09), Apr 13-16, 2009, San Francisco, USA.

M. García-Valls, A. Alonso, and J. A. de la Puente. Dynamic Adaptation Mechanisms in Multimedia Embedded Systems. In Proc. of the 7th IEEE International Conference on Industrial Informatics. (IEEE INDIN 09). Cardiff, UK. 24-26 June 2009.

M. García-Valls, A. Alonso, and J. A. de la Puente. Mode Change Protocols for Predictable Contract-Based Resource Management in Embedded Multimedia Systems. In Proc. of the 6th IEEE International Conference on Embedded Software and Systems. (IEEE ICESS 09). Digital

Object Identifier 10.1109/ICESS.2009.93, pp. 221-230. HangZhou, Zhejiang, China. 25-27 May 2009.

M. García-Valls, A. Alonso, and J. A. de la Puente. Dynamic Adaptation Mechanisms in Multimedia Embedded Systems. In Proc. of the 7th IEEE International Conference on Industrial Informatics. (IEEE INDIN 09). Cardiff, UK. 24-26 June 2009.

M. García-Valls, A. Alonso, and J. A. de la Puente. Time-Predictable Reconfiguration with Contract-Based Resource Management. In Proc. Of the 4th International IEEE Workshop on Service Oriented Architectures in Converging Networked Environments. (IEEE SOCNE'09) Bradford, UK. 26 May 2009.

P. Basanta-Val, I. Estévez Ayres, M. García-Valls, and Luis Almeida. A synchronous scheduling service for distributed real-time Java.IEEE Trans. on Parallel and Distributed Systems. August 2009.

Pedro Silva, Luis Almeida, Daniele Caprini, Tullio Facchinetti, Francesco Benzi, and Thomas Nolte, "Experiments on timing aspects of DC-Powerline communications", Proceedings of the 14th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA), Palma De Mallorca, Spain, 22-26, September, 2009.

Ezio Bassi, Francesco Benzi, Tullio Facchinetti, Luis Almeida, Thomas Nolte, "Powerline Communication in Electric Vehicles", Proceedings of the IEEE International Electric Machines and Drives Conference (IEMDC), Miami, Florida, USA, 3-6, May, 2009.

Javier Silvestre, Ricardo Marau, Paulo Pedreiras, Luis Almeida. On-line QoS Management for Multimedia Real-Time Transmission in Industrial Networks, IEEE Transactions on Industrial Electronics (to appear).

R. Santos, A. Vieira, R. Marau, P. Pedreiras, A. Oliveira, L. Almeida, T. Nolte. Implementing Server-Based Communication within Ethernet Switches. CRTS 2009 – 2nd Workshop on Compositional Theory and Technology for Real-Time Embedded Systems (satellite of RTSS 2009), Washington DC, USA, December 1, 2009.

S, Samii, P. Eles, Z. Peng and A. Cervin, Quality-Driven Synthesis of Embedded Multi-Mode Control Systems, Proc. 46th Design Automation Conference (DAC), San Francisco, CA, 2009.

S. Samii and A. Cervin, P. Eles, Z. Peng, Integrated Scheduling and Synthesis of Control Applications on Distributed Embedded Systems, Proc. Design, Automation & Test in Europe (DATE), 2009.

Narasinga Rao Miniskar, Elena Hammari, Satyakiran Munaga, Stylianos Mamagkakis, Per Gunnar Kjeldsberg, Francky Catthoor: Scenario Based Mapping of Dynamic Applications on MPSoC: A 3D Graphics Case Study. In Embedded Computer Systems: Architectures, Modeling, and Simulation, 9th International Workshop, SAMOS 2009, Samos, Greece, July 20-23, (2009)

Bartzas, A.; Peon, M.; Mamagkakis, S.; Catthoor, F.; Soudris, D. and Mendias, J.: Direct memory access usage optimization in network applications for reduced memory latency and energy consumption. In Journal of Embedded Computing, p.241-254, Vol.3, Issue 3, (2009)

***The above are new references, not present in the Y1 deliverable.***

## *2.5 Keynotes, Workshops, Tutorials*

### Keynotes

Alan Burns,
Adaptive/reconfigurable servers,
Compositional Theory and Technology for Real-Time Embedded Systems (CRTS,2009), Washington, USA, December 2nd, 2009.

Stylianos Mamagkakis
Emerging multicore hardware platforms and their software support challenges,
21st EUROMICRO Conference on Real-Time Systems - ECRTS 2009,
Dublin, Ireland, July 2nd, 2009
http://ecrts09.dsg.cs.tcd.ie/keynote-speaker.php

Stylianos Mamagkakis
Adaptive solutions for the emerging reliability and multicore resource management challenges,
2nd Workshop on Adaptive and Reconfigurable Embedded Systems - APRES 2009,
Grenoble, France, October 11th, 2009 within ESWEEK 2009
http://www.artist-embedded.org/artist/Keynote.html

Marisol García-Valls
Software models for distributed real-time data distribution in avionics,
Master on Aircraft Systems Integration, EADS-UC3M
Airbus center, Getafe (Madrid), October 28th, 2009.

Marisol García-Valls
Resource Management for Multimedia Embedded Systems,
Master on Industrial Automation and Informatics, Department on Systems Engineering and Automation,
Technical University of Valencia, Spain, July 1st-3rd 2009.

Luis Almeida
A Dynamic Scheduling Approach to Designing Flexible Safety-Critical Systems
Seminar at the CISTER research Unit, Polytechnic Institute of Porto, Portugal, 5 June 2009.

### Workshops

OSPERT 2009
http://www.artist-embedded.org/artist/-OSPERT-2009-.html
Fifth International Workshop on Operating Systems Platforms for Embedded Real-Time Applications, Dublin, Ireland, June 30, 2009.

XII Spanish Workshop on Real-Time Systems

http://www.it.uc3m.es/jtr2009

University Carlos III de Madrid, 5-6 February, 2009.

CRTS 2009

http://www.artist-embedded.org/artist/CRTS-2009.html

Second Workshop on Compositional Theory and Technology for Real-Time Embedded Systems, Washington DC, USA, December 1, 2009.

IRTAW 2009

14th International Real-Time Ada Workshop, Italy, 7-9 September (sponcered in part by ArtistDesign).

## Tutorials

Alan Burns

Scheduling and Timing Analysis for Safety Critical Systems

Safety Critical Systems Club (SCSC), London, Oct 2009.

Tutorial focused on verification of the timing properties of safety-critical real-time systems, looking at the standard ways of scheduling and analysing the concurrent software that typically executes within embedded systems. These techniques are crucial for safety-critical systems where there is a need for a high level of assurance that timing constraints will be satisfied in all situations, even those that are very rare.

http://www.safety-club.org.uk/main.html

***The above is new material, not present in the Y1 deliverable.***

# 3. Milestones, and Future Evolution

## *3.1 Problems to be Tackled over the next 12 months (Jan 2010 – Dec 2010)*

The main cluster-wide problem to be tackled during the next 12 months is to continue to develop the wiki addressing the taxonomy of resource usage that has 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 will be covered. A broad range of resources will be included – see current wiki. Work has already started on multi-resource modelling techniques, this will continue.

In addition, all the 17 technical activities described in Section 2 will continue, as will the work addressing the millstones described below.

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

***Essentially the same objectives to those contained in the Year 1 deliverable.***

## *3.2 Current and Future Milestones*

We continue to work on the milestones set during Year 1:

- 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.
- Extend sensitivity analysis to EDF and 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 addition in Year 3 we will consider:

- Optimizing reservation parameters for partitioning parallel real-time applications to multi-core platforms.
- Determining preemption points to reduce cache-related preemption delays in task execution.
- Development of the first bin-packing based non-migrative real-time scheduling algorithm for a heterogeneous multicore processor.
- Prototype implementation of the cooperative service execution for resource constraint devices.

### *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 Year2 this activity has benefited from direct funding from

- JEOPARD – EU funded project on parallel real-time Java, consortium includes Vienna and York.
- eMuCo – EU funded project aimed at developing a platform for future mobile devices based on multi-core architecture, consortium included York, the Ruhr-Universität Bochum, Technische Universität Dresden and Politehnica University of Timisoara, as well as industrial partners, such as Infineon, Telelogic, ARM and GWT-TUD.
- Reflect - Reflection Mechanisms in Real-Time Embedded Systems – Portuguese funded FCT project involving the Polytechnic Institute of Porto.
- CooperatES - QoS-Aware Cooperative Embedded Systems – Portuguese funded FCT project involving the Polytechnic Institute of Porto.
- RESCUE - REliable and Safe Code execUtion for Embedded systems – Portuguese funded FCT project involving the Polytechnic Institute of Porto.
- The FRESCOR EU IST project, in which the following ArtistDesign partners were involved: University of Cantabria, University of York, Scuola Superiore Santa Anna, Technical University of Kairserslautern, Technical University of Valencia, see www.frescor.org.
- RT-MODEL Spanish project, in which the following ArtistDesign partners are involved: Technical University of Madrid, University of Cantabria,Technical University of Valencia
- ASSERT - Automated proof based System and Software Engineering for Real-Time applications. EU funded project. Main objective is to improve the system-and-software development process for critical embedded real-time systems in the Aerospace and Transportation domains.
- MADEJA: French funded (Region Rhône-Alpes) in which VERIMAG was involved for studying quantitative prediction of dynamic memory allocation for RTSJ scoped-memory mechanism.
- 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/
- RESCORE. Portuguese funded project aiming to develop real-time scheduling algorithms for multicores.
- HARTES - HArd Real-Time Ethernet Switching - Portuguese funded FCT project involving the University of Aveiro.
- OMP IST STREP – Incorporating QoS and Resource Management APIs in Open Media Platform standard OpenMax.

- 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.
- Tempo – Nationally funded UK project at York focusing on multicore platforms and timing analysis/scheduling.
- National Project on "Wireless Sensor Networks for Monitoring Patients with Heart Deseases"ArtistDesign partners involved: Pisa, Evidence.
- REHEAT – Nationally funded project at Porto focusing on parallel processing platforms http://www.cister.isep.ipp.pt/activities/REHEAT.
- Open Media Platform, an EU funded project looking at resource management; involves NXP.
- INDEXYS, INDustrial EXploitation of the genesYS cross-domain architecture; involves Vienna, NXP, Kaiserslautern.
- 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
- CHAT - EU STREP on networked automation systems involving ULUND plus several non-ARTIST partners
- Flexibility for TTEthernet, Kaiserslautern.
- iLAND – mIddLewAre for deterministic dynamically reconfigurable NetworkeD embedded systems – ARTEMIS funded project in which the following partners are involved: University Carlos III of Madrid and University of Porto.

***In the above list the final ten projects started in 2009.***

# 4. Internal Reviewers for this Deliverable

- **Michael González Harbour** (Universidad de Cantabria)
- **Peter Puschner** (Technische Universitaet Wien)
- **Andy Wellings** (University of York)