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

Activity Progress Report for Year 4

Resource-aware Operating Systems

Cluster:

Operating Systems and Networks

Activity Leader:

Prof. Giorgio Buttazzo (Scuola Superiore S. Anna)
http://retis.sssup.it/~giorgio/

Policy Objective (abstract)

The objective of this activity is to show how current operating systems can be designed to support emerging real-time applications that exhibit a high degree of complexity and operate in highly dynamic environments where resource demands can change unpredictably. Adaptive resource management mechanisms will be investigated both in uniprocessor and multicore architectures. The impact on operating system standards (like RT-POSIX and OSEK) will also be taken into account.

Versions

number	comment	date
1.0	First version delivered to the reviewers	January 30th 2012

Table of Contents

1. Overview of the Activity 3
1.1 ArtistDesign participants and their role within the Activity 3
1.2 Affiliated participants and their role within the Activity 3
1.3 Starting Date, and Expected Ending Date 4
1.4 Policy Objective 4
1.5 Background 5
1.6 Technical Description: Joint Research 6

2. Work Achieved in the NoE 8
2.1 Synthesis View of the Main Overall Achievements 8
2.2 Work achieved in Year 1 (Jan-Dec 2008) 9
2.3 Work achieved in Year 2 (Jan-Dec 2009) 10
2.4 Work achieved in Year 3 (Jan-Dec 2010) 10
2.5 Work achieved in Year 4 (Jan-Dec 2011) 11

3. Detailed view of the progress in Year 4 *(Jan-Dec 2011)* 13
3.1 Technical Achievements 13
3.2 Individual Publications Resulting from these Achievements 16
3.3 Interaction and Building Excellence between Partners 20
3.4 Joint Publications Resulting from these Achievements 20
3.5 Keynotes, Workshops, Tutorials 23

4. Internal Reviewers for this Deliverable 28

1. Overview of the Activity

1.1 ArtistDesign participants and their role within the Activity

Cluster Leader: Giorgio Buttazzo – Scuola Superiore S. Anna (Italy)
Role: Activity coordinator, kernel maintenance, development of robotic applications.

Team Leader: Luis Almeida – University of Aveiro (Portugal)
Role: networking platform, development of distributed applications.

Team Leader: Gerhard Fohler – University of Kaiserslauten (Germany)
Role: video streaming applications, scheduling.

Team Leader: Michael Gonzalez Harbour – University of Cantabria (Spain)
Role: definition of the POSIX operating system interface.

Team Leader: Alan Burns – University of York (UK)
Role: feasibility analysis of fixed priority real-time systems.

Team Leader: Eduardo Tovar – Polytechnic Institute of Porto (Portugal)
Role: distributed applications and QoS over heterogeneous networks.

-- Changes wrt Y3 deliverable --

No changes with respect to Year 3.

1.2 Affiliated participants and their role within the Activity

Team Leader: Tullio Facchinetti – University of Pavia (Italy)
Role: embedded real-time systems and robotics applications.

Team Leader: Paolo Gai – Evidence s.r.l. (Italy)
Role: real-time kernels and operating systems standards.

Team Leader: Hermann Haertig – Universtity of Dresden (Germany)
Role: microkernel architectures and virtualization techniques

Team Leader: Sverre Hendseth (Norwegian University of Science and Technology (Norwey)
Role: distributed systems and concurrent languages.

Team Leader: Pau Marti – Universitat Politècnica de Catalunya (Italy)
Role: control applications and schedulability of event-driven control systems.

Team Leader: Alejandro Alonso – Technical University of Madrid (Spain)
Role: QoS resource management, high integrity systems

Team Leader: Marisol García Valls – University Carlos III of Madrid (Spain)
Role: QoS-based resource management, real-time support in middleware and distributed systems.

Team Leader: Alfons Crespo –Technical University of Valencia (Spain)
Role: real-time memory management, virtualization of real-time kernels

Team Leader: Salvatore Scafidi – Windriver (Italy)
Role: real-time operating systems

Team Leader: Stylianos Mamagkakis – IMEC Leuven (Switzerland)
Role: Energy-aware scheduling and memory management for dynamic applications

Team Leader: Marko Bertogna – University of Modena (Italy)
Role: real-time operating systems

Team Leader: Antonio Bersani – Microchip Technology (Italy)
Role: Hardware components and I/O devices for real-time applications

-- Changes wrt Y3 deliverable --

Norwegian University of Science and Technology (Norwey) and University of Modena (Italy) have been added as affiliated partners.

1.3 Starting Date, and Expected Ending Date

Starting date: January 1st, 2008

Ending date. December 31st, 2011.

-- Changes wrt Y3 deliverable --

No change.

1.4 Policy Objective

The main objective of this activity is to investigate how current real-time operating systems have to be extended or modified to support emerging real-time embedded systems characterized by a high degree of complexity, highly variable resource requirements and parallel processing such as multicores. Most embedded systems are often characterized by scarce resources, in terms of processing power, memory, space, weight, energy, and cost. Hence, another objective is to investigate kernel mechanisms that can efficiently manage the available resources, taking multiple constraints into account, whilst guaranteeing isolation properties. Also, to support dynamic applications with variable resource requirements or to cope with unpredictable resource availability, feedback control techniques for resource management at the operating system and application level will be investigated. The impact on operating system standards (like RT-POSIX and OSEK) will also be taken into account. In fact, developing real-time applications and components using an interface compliant to a standard will promote portability to other compliant platforms and will challenge the current standard to be extended to better meet the needs of advanced applications with flexible resource requirements. We realize though that in specific application domains, significant performance advantages can be realized by optimizing software across layers, for example exploiting specific behaviour of a medium access control protocol. This is often the case in operating systems for sensor network platforms such as TinyOS or NanoRK. Such cross-layer design does not necessary contradicts operating system standards, but they do require other interfaces.

-- Changes wrt Y3 deliverable --

No changes with respect to Year 3.

1.5 Background

Although there is a large variety of real-time operating systems (RTOSs) varying in sizes, level of provided services, and efficiency, there are some common elements that can be found in most of them:

- An RTOS usually provides support for concurrent programming via processes or threads or both. Processes usually provide protection through separate address spaces, while threads can cooperate more easily by sharing the same address space, but with no protection.
- Real-time scheduling services are provided because this is one of the keys to obtaining a predictable timing behaviour. Most current RTOS's provide the notion of a scheduling priority, usually fixed, as for the moment there are few systems providing deadline-driven or other dynamic-priority scheduling.
- Although some RTOS designed for high-integrity applications use non preemptive scheduling, most support preemption because it leads to smaller latencies and a higher degree of utilization of the resources.
- The OS has to support predictable synchronization mechanisms, both for events or signal and wait services, as well as for mutual exclusion. In the later case some way of preventing priority inversion is required because otherwise very improbable but also very long delays may occur.
- The OS has to provide time management services with sufficient precision and resolution to make it possible for the application to meet its timing requirements.
- OS behaviour should be predictable, and so metrics of the response time bounds of the services that are used in real-time loops should be clearly given by the RTOS manufacturer or obtained by the application developer. These metrics include the interrupt latency (i.e., time from interrupt to task run), the worst case execution time of the system calls used in real-time loops, and the maximum time during which interrupts are masked or disabled by the OS and by any driver.

An RTOS is generally chosen not only for its real-time characteristics, but also for the middleware that is integrated in the RTOS, such as file system, communication stack, for its portability to different platforms (i.e., the board support packages that are provided), and for the associated cross-development environment.
A commercial RTOS is usually marketed as the run-time component of an embedded development platform, which also includes a comprehensive suite of (cross-) development tools and utilities and a range of communications options for the target connection to the host, in an Integrated Development Environment (IDE). Moreover, the vendor generally provides development support. For each successful open source RTOS there is also at least one commercial distributor that provides development tools and development support. For many embedded-systems companies, the availability of development tools and support is a major requirement for choosing a particular RTOS. The quality of the overall package deal, including service and pricing strategy is often decisive in choosing a particular RTOS.

-- Changes wrt Y3 deliverable --

No changes with respect to Year 3.

1.6 Technical Description: Joint Research

Preemption Avoidance in Fixed Priority Scheduling
Preemptions in computing systems lead to extended response times of a task, not only due to the actual time of a higher priority task executing, but also due to the loss of working set in caches and a like, which needs to be reloaded during execution. When compared to other approaches regular fixed-priority scheduling suffers from a larger number of preemptions. On the other hand, it also offers lower complexity scheduling, when compared to other approaches like EDF. To alleviate this, University of Porto, Pisa and Modena porposed different approaches for avoiding/collating preemptions, by enforcing non-preemptible regions. Of particular concern is that opposed to related work, the approach focusses on floating non-preemptive regions, which do not required modifications in the application code.

Heterogenous Multicore Scheduling
Heterogenous Multicores are one logic next step after symmetric multicore processors. Porto has set out in late 2010 to investigate the use of such architectures in real-time systems. Of particular concern in this is resource sharing. The approach developed considers a two-type heterogenous platform and a number of shared resources, which are protected via non-preemtible critical sections. The approach limits the speedup required when comparing with an optimal scheduling approach.

Multicore embedded real-time systems
Next generation RTOS must allow optimal off-line partitioning of the application source code on the different CPUs available on multicore heterogeneous systems, as well as on-line strategies for the run-time migration with the objective of guaranteeing optimal usage of the CPUs available, with real-time response as well as minimization of power consumption. The Cluster investigated a method for partitioning parallel applications on top of multicore platforms, optimizing bandwidth usage and tacking timing and precedence constratins into account.

Component-based operating systems
To optimize the use of resources and increase software portability on different platforms, it is highly desirable to compose the operating system using the functions strictly necessary for the application. To achieve this goal, it is crucial to design the operating system to be modular, so that each component can be independently developed from the others and can be replaced without changing the application. The Cluster was involved in the development of modular real-time systems and hierarchical scheduling techniques that allow isolating the temporal behavior of real-time components.

Deadline Scheduling on Linux
Kernel support for deadline-based scheduling algorithms has been developed in Linux. The contribution to the cluster has been twofold: On one hand, a new POSIX-like scheduling policy, called SCHED_DEADLINE, was implemented as a novel scheduling class, like the POSIX-specified ones already present in the kernel; on the other hand, the existing bandwidth control mechanism (called throttling), was extended using a Constant Bandwidth Server.

Real-Time techniques for Service oriented Architectures (SoA)
Resource provisioning in SoA is a challenge, especially when the system consists of hundreds of nodes and real-time constraints are required. We extended a technique called “Advance Reservation” that is used in GRID architectures to cope with QoS requirements and dynamic admission control which is required by very dynamic SoA. The Cluster adopted EDF and resource reservations as a base scheduler. On top of it, two admission control methods have been implemented: a deterministic admission control, which is safe but can waste resources, and a probabilistic admission control, which greatly improves on resource utilization.

-- Changes wrt Y3 deliverable --

More effort has been put to multiprocessor systems, heterogeneity, and abstraction techniques for component-based design.

2. Work Achieved in the NoE

2.1 Synthesis View of the Main Overall Achievements

The main achievents of the cluster in the area of real-time operating systems can be summarized as follows:

Operating systems support for multicore platforms. A joint work between Pisa, Lund, TUKL, and Evidence resulted in developing a programming framework for supporting parallel applications on multicore platforms, providing resource reservation, real-time guarantees, temporal isolation, and resource optimization algorithms. The resulting approach has been considered by Ericsson for software development in next generation cell phones.

Deadline scheduler in the Linux kernel. An efficient real-time scheduler (called SCHED_DEADLINE), based on the Earlieast Deadline First algorithms with a support for resource reservation, has been implemented in the Linux kernel and considered by the Linux community to be integrated in the main line of the kernel.

Integration of real-time scheduling with cache-aware analysis. A tight collaboration with the Cluster on Compilers and Timing Analysis resulted in the integration of cache analysis with limited preemptive scheduling, as a method for improving the estimation of worst-case execution times, reducing its variability and enhancing the schedulability level of fixed priority systems. The work has been carried out in collaboration with Airbus (for avionic applications) and Bosch (for automotive systems).

Power-aware resource management. A Power Manager has been developed on the ERIKA real-time kernel to manage power and temperature constraints in real-time embedded systems consisting of CPU and multiple devices working in different energy modes. The Power Manager has been developed as an open source layer for providing a uniform and portable approach for energy management.

Cycle accurate simulation tool for real-time applications on multicore platforms. A tight collaboration with the Cluster on Execution Platforms allowed integrating the Erika real-time kernel on the MPARM simulator developed by the University of Bologna. As a result, real-time source code can be executed on a simulated multi-core platform for testing the application performance on custom scheduling algorithms, cache policies, and energy-aware management algorithms.

Altogether, these contributions are reported in 190 publications produced along the 4 years of the activity and distributed as follows:

Year	Individual publications	Joint publications
Year 1	15	14
Year 2	19	17
Year 3	38	16
Year 4	43	28
Total	**115**	**75**

Moreover, the partners involved in this activity also participated in several European projects, some of which have already been referred before:

- FRESCOR, ACTORS, PREDATOR, IRMOS, SO(O)S.

In this aspect, we believe that the activity played a significant role in promoting project submissions and cross-fertilization of ideas and results. This cross-fertilization was also fostered by several workshops on related topics organized with a strong involvement from activity partners. Some of such workshops were consistently organized along the years and others occurred on an event basis, as well as tracks and special sessions organized in related conferences, essentially in the areas of Real-Time Systems, Embedded Systems and Industrial Automation:

Finally, the activity also carried out a considerable dissemination effort. In fact, activity members participated in almost all summer schools organized by ArtistDesign, either in Europe, China, South America and Morocco.

-- The above is new material, not present in the Y3 deliverable --

2.2 Work achieved in Year 1 *(Jan-Dec 2008)*

- **Pisa**: Partitioning a real-time application in a multi-core architecture, with the objective of providing the resource reservation abstraction to achieve temporal protection and to allocate a fraction of the available resources to a given application. Unfortunately, extending the resource reservation paradigm to multicore architectures is not trivial, since resource allocation must be considered together with the problems of exploiting the intrinsic parallelism of certain applications.
- **Pisa**: Another issue we started to investigate in this first year is the evaluation of the interactions of cache memories with scheduling. In particular, preemptive algorithms tend to destroy the cache content of the preempted activity, so increasing the number of cache misses, causing an increase of task computation times. To better evaluate such dependencies we decided to perform a set of simulation experiments aimed at measuring such effects in different scenarios and working conditions.
- **Pavia**: We started investigating the possibility of executing different operating systems on a single multicore platform, by assigning a different core to each operating system, as a special kind of virtualization technique that does not rely on lower-layer components. The short-term goal is to run a real-time operating system along with a general-purpose one in a completely transparent fashion, while the subsequent research will address the cross effects between the operating systems due to shared resources, like memory, bus, peripherals and cache.
- **Pavia**: Specific work on embedded and robotics applications has been started. In this domain the research directions focuses on the development of applications based on small micro-controllers to assess the benefits of real-time computing for the predictability and the overall performance and the application.
- **Catalonia**: We also started to investigate the schedulability analysis of a specific type of control applications named event-driven control applications. The importance of these type of controllers is that they can provide the same control performance than standard periodic controllers while minimizing resource utilization.

-- No changes wrt Y3 deliverable --

This section was already presented in the Y3 deliverable, in section 1.7.

2.3 Work achieved in Year 2 *(Jan-Dec 2009)*

Resource reservations on multi-core platforms
Extending the resource reservation framework from uniprocessor to multicore systems is not trivial and requires new theoretical bases. In this year, the problem of partitioning parallel real-time applications on top of multicore platforms has been investigated for exploiting the available parallelism offered by modern architectures.

Deadline Scheduling on Linux
Extending the Linux kernel to support novel real-time scheduling algorithms is essential to expose the most relevant research results achieved within small research communities to the entire world of software developers. A great effort has been done in this Cluster to introduce the Earliest Deadline First (EDF) scheduler as a scheduling class in Linux, to allow a more efficient support of real-time applications and enable the implementation of advanced resource reservations techniques.

Power-Aware Library
The Scuola Superiore Sant'Anna of Pisa started a work to implement a kernel library to manage power and temperature constraints in real-time embedded systems.

ERIKA on MPARM
The Scuola Superiore Sant'Anna of Pisa ported the Erika real-time kernel on the MPARM simulator developed by the University of Bologna.

-- No changes wrt Y3 deliverable --

This section was already presented in the Y3 deliverable, in section 1.8.

2.4 Work achieved in Year 3 *(Jan-Dec 2010)*

Resource reservations on multi-core platforms
The work on multicore resource reservation has been extended by providing a methodology that can partition a parallel application on a set of uniprocessor reservations with minimal bandwidth requirements.

Deadline Scheduling on Linux
A great effort has been done in this Cluster to introduce resource reservation and deadline-based scheduling (EDF) in the Linux operating system. The implementation is now complete and allows a more efficient support of real-time applications, enabling the implementation of advanced resource reservations techniques with adaptive bandwidth management.

Power-Aware Resource Management
The work started by the Scuola Superiore Sant'Anna of Pisa on providing kernel support for managing power and temperature constraints has been completed. A C library for the Erika kernel has been developed to manage real-time applications with power and temperature requirements. A uniform approach has been adopted to handle different types of resources, such as processors, memories, buses, wireless transceivers, and I/O devices.

Architecture simulator for real-time applications
The multi-processor cycle-accurate architectural simulator (MPARM) developed at the University of Bologna has been extended to run real-time concurrent applications on the Erika kernel developed by Evidence and the Scuola Superiore Sant'Anna of Pisa. The simulator has been used to perform a set of experiments of non-preempitive scheduling, cache-aware scheduling, and power-aware strategies.

-- No changes wrt Y3 deliverable --

This section was already presented in the Y3 deliverable, in sections 1.8 and 3.1.

2.5 Work achieved in Year 4 *(Jan-Dec 2011)*

UNIBO-PISA started a collaboration on resource optimization and scheduling in multiprocessor platforms. The interaction has been carried out by integrating methods and tools developed by each partner, such as the Erika RTOS and the elastic method developed by SSSA, and the MPARM and the optimization algorithms developed by UNIBO.

PISA-UC3M initiated a research collaboration on resource management policies. A PhD student from PISA made a research stay during the academic year 2010/2011 at UC3M.

USAAR, PISA, Dortmund, AbsInt. Also supported be the PREDATOR project, these partners continued the collaboration to improve the estimation of worst-case execution times by integrating cache-aware compilers with limited preemptive scheduling.

EVIDENCE-PISA. A great effort has been done to introduce resource reservation and deadline-based scheduling (EDF) in the Linux operating system, so enabling the implementation of advanced resource reservations techniques.

LUND-TUKL-PISA. Also supported by the ACTORS project, these partners collaborated to develop a design framework for partitioning real-time applications on multicore heterogeneous systems, with the objective of guaranteeing optimal usage of the available resources. A graphic tool has been also developed for testing the algorithms on different real-time applications.

York, Cantabria, Porto, Madrid, Valencia collaborated for providing a language support for programming schedulable systems. 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.

Cantabria, Madrid: UPM and UC3M, Bilbao, UnivPorto collaborated for providing real-time support to middleware and composability. A set of timing analysis tools has been integrated with a toolset for MDE. In addition, a new approach has been explored to integrate the real-time end-to-end flow model with the automatic generation of Ravenscar-compliant source code in distribution middleware.

Madrid, Pisa, Aveiro, UnivPorto, Malardalen, NXP, TUKL worked on protocol optimizations for embedded real-time communications. The validation showed performance improvements in comparison to currently used infrastructures. The performance has been reported to a journal in an article which now in accepted status.

Catania, Pisa, Evidence have been involved in intelligent transportation systems, automatic traffic monitoring and road surveillance. Various sensors have been used to estimate traffic parameters. Catania proposed a wireless sensor network architecture based on computer vision techniques for automatic scene analysis and interpretation.

Cantabria, Padova. The University of Cantabria and the University of Padova have been working on methods for the design of real-time component-based applications on open platforms. Techniques are being developed to incorporate schedulability analysis to the component models. A student exchange was derived from this collaboration.

Cantabria, Madrid (UPM) have developed an approach based on the Model-Driven Engineering (MDE) paradigm to ease the automatic generation of HDRT applications from high-level system models. The real-time end-to-end flow model has been integrated with with the automatic generation of Ravenscar-compliant source code in distribution middleware.

Madrid (UC3M), Madrid (UPM) have developed a scheduling mechanism for dynamic priority assignment in QoS-based resource management based on contracts and budget scheduling. It has been applied to embedded multimedia systems in a way that safe budget overruns are allowed to increase the overall delivered quality.

-- The above is new material, not present in the Y3 deliverable --

3. Detailed view of the progress in Year 4 *(Jan-Dec 2011)*

3.1 Technical Achievements

Partitioning real-time applications on multi-core platforms (Pisa)
The RETIS group of Pisa investigated a method for partitioning parallel applications on top of multicore platforms, optimizing bandwidth usage and taking timing and precedence constraints into account [PISA_1]. The approach has been implemented into a graphic interactive tool that allows the user to define a parallel application with precedence relations and partition it on a set of reservations. The tool is available from: http://particore.sssup.it

Power management in real-time embedded systems (Pisa)
The RETIS Lab of Pisa continued to work on energy management from different perspectives. From one hand, a kernel support library for the Erika kernel has been developed to manage real-time applications with power and temperature requirements [PISA_2]. A uniform approach has been adopted to handle different types of resources, such as processors, memories, buses, wireless transceivers, and I/O devices. The approach has been adopted to save energy in body sensor networks for mobile health monitoring systems [PISA_3]. On the other hand, a new approach has been proposed to save energy in distributed systems with execution, communication, and redundancy requirements [PISA_4].

Multicore and manycore real-time scheduling (Pisa)
The ReTIS lab of Pisa has pursued two research lines: efficient multicore scheduling in the Linux kernel; and real-time scheduling in manycore Networks on Chip. For the first topic, we continued to develop the SCHED_DEADLINE scheduler for the Linux kernel [PISA_5], [PISA_6]. We improved over our previous realisation by implementing a in-kernel shared data-structure that speeds up the scheduling decision, and achieves good scalability with respect to the number of processors. The result has been published at the OSPERT workshop in July 2011 [PISA_7]. We also investigated many-core architectures, that are considered to be the future of multicore chip. In particular we performed several experiments on the Intel SCC chip which consists of 48 cores connected by a mesh network-on-chip [PISA_8]. We measured the communication cost using different shared memory mechanisms, both using local memory, non-coherent caches and global memory. Results demonstrate that local memory and shared caches show the best performance. In the same topic, we have proposed methodology for allocation of task pipelines on multicore architectures.

Core to core communications and deployment in tile-based architectures (Pisa)
Within the S(o)OS project (http://www.soos-project.eu/), the Pisa group is investigating novel Operating System architectures and components supporting scheduling and deployment for massively parallel and distributed systems of tomorrow (e.g., tera/exa-scale many-core platforms). Prelilminary experiments have been conducted on the Single-Chip Cloud Computer (Intel SCC), featuring a 48-core tile-based Network-on-a-Chip hardware architecture, without memory coherency, where consistency of the OS is guaranteed by running a distributed version of Linux, i.e., 48 instances of the Linux OS and kernel (configured without any multi-core support) running independently from one another over the 48 cores. Different core to core communication mechanisms have been compared, demonstrating that, depending on the actual workload, the fast core to core message-passing support provided by Intel to the user-space applications may not always be the best solution [PISA_14] [PISA_15] [PISA_16]. Rather, in certain scenarios, the globally shared memory mechanism, coupled with proper cache flushing and invalidation operations (needed to ensure data consistency across multiple cores) shows a better performance.

Real-Time Multimedia Applications on Service-oriented Infrastructures (Pisa)
The KVM hypervisor was added under Linux, integrating real-time virtualization capabilities into the ISONI (Intelligent Service-Oriented Networking Infrastructure) IaaS infrastructure, realized within the IRMOS project (http://www.irmosproject.eu/). Specifically, a hard-reservation variant of the Constant Bandwidth Server (CBS) algorithm is used to achieve temporal isolation across threads or processes [PISA_9] [PISA_10] [PISA_11]. The scheduler is configured through the use of the cgroups-based user-space/kernel-space interface, based on the Linux virtual filesystem (e.g., creating a reservation is done by creating a folder in the virtual file-system, and budget and period is assigned by special entries in such a folder). An inner scheduler implements Fixed Priority scheduling across the threads attached to the reservation [PISA_12] [PISA_13]. The scheduler also supports multi-processor platforms.

Preemption Avoidance in Fixed Priority Scheduling (Porto-IPP)
Preemptions in computing systems lead to extended response times of a task, not only due to the actual time of a higher priority task executing, but also due to the loss of working set in caches and a like, which needs to be reloaded during execution [IPP_1] [IPP_2]. When compared to other approaches regular fixed-priority scheduling suffers from a larger number of preemptions. On the other hand, it also offers lower complexity scheduling, when compared to other approaches like EDF. To alleviate this, Porto has developed an approach, which avoids/collates preemptions, by enforcing non-preemptible regions. Of particular concern is that opposed to related work, the approach focusses on floating non-preemptive regions, which do not required modifications in the application code.

Application areas are fixed priority environments where non-preemptive scheduling is not possible due to greatly varying rates of applications. Examples include the integration of previously isolated functionality on to a single more powerful device as can be observed in the automotive domain.

Heterogenous Multicore Scheduling (Porto-IPP)
Heterogenous Multicores are one logic next step after symmetric multicore processors [IPP_3] [IPP_4]. Porto has set out in late 2010 to investigate the use of such architectures in real-time systems. Of particular concern in this is resource sharing. The approach developed considers a two-type heterogenous platform and a number of shared resources, which are protected via non-preemtible critical sections. The approach limits the speedup required when comparing with an optimal scheduling approach.

Provably Good Resource Sharing for Multiprocessors (Porto-IPP)
The number of processor cores in multicore processor on a single chip is increasing exponentially and this brings a tremendous amount of raw computing capacity available to designers. Many applications however are composed of tasks which share data or other resources and this can limit the parallelism in the application and consequently making it impossible to harness the full power of a multicore processor. The research literature makes available several protocols for allowing real-time tasks executing on a multiprocessor to share resources but unfortunately, it has been difficult to know how well they perform and in particular, it has been difficult to know if there is any task set for which existing resource sharing protocols perform very poorly (for example performing like a uniprocessor where it would be possible to utilize the entire parallism in the computing platform).
Researchers at CISTER research unit have developed a new resource sharing protocol for multiprocessors and a scheduling algorithm to be used in conjunction with this protocol [IPP_5]. This protocol has provably good performance, meaning that if it is possible for a set of tasks to meet deadlines then tasks meet deadlines as well if this resource sharing protocol is used (and its corresponding processor scheduling algorithm) assuming that the new protocol is given processors which are speeded up by a constant factor.

This may be applicable in the automotive, industrial electronics or space domains, all of which see an unprecedented deployment of multicores.
URL: http://www.cister.isep.ipp.pt/docs/

Task assignment for heterogeneous multiprocessors (Porto-IPP)
Multiprocessors implemented on a single chip (a.k.a multicores) are becoming ubiquitous. More and more designers realize the advantages (in terms of computational speed or power "consumption") in using multicore processors comprising different functional units or processor units specialized to operate at very high speed for the type of computation it is designed for. For example a graphics processor unit can perform graphics operations at a much higher speed than a general purpose processor. From the perspective of real-time scheduling, this implies that the execution time of a task depends on which processor core it is assigned to. Previously known algorithms have high computational complexity but one can observe that many heterogeneous multiprocessors used in practice have only two types of processors.
Researchers at CISTER research unit at ISEP/IPP have therefore considered the special case of a heterogeneous multiprocessor with two types of processor and under this assumption they have developed a new algorithm for assigning real-time tasks to processors. What makes this new algorithm remarkable is that it offers a so-called resource augmentation bound similar to the prior state of art but in addition to that, in experimental performance evaluations, it is found to be able to succeed to assign more task sets than prior state of art and it runs more than 10000 times faster [IPP_6].

Again this may find applicability in different domains. In this case, we are looking more towards multimedia type platforms, which support different specialized hardware processors which are still general enough to be exploited for different media types.
URL: http://www.cister.isep.ipp.pt/docs/

Combining Multicore Communications API with Compile-Time Virtualisation (York)
Within the domain of embedded systems, hardware architectures are commonly characterised by application-specific heterogeneity. Systems may contain multiple dissimilar processing elements, non-standard memory architectures, and custom hardware elements. The programming of such systems is a considerable challenge, not only because of the need to exploit large degrees of parallelism but also because hardware architectures change from system to system. To solve this problem, the York group proposes a novel combination of a new industry standard for communication across multicore architectures (MCAPI), with a minimal-overhead technique for targeting complex architectures with standard programming languages (Compile-Time Virtualisation) [YORK_1]. The Multicore Association have proposed MCAPI as an industry standard for on-chip communications. MCAPI abstracts the on-chip physical communication to provide the application with logical point-to-point unidirectional channels between nodes (software thread, hardware core, etc.). Compile-Time Virtualisation is used to provide an extremely lightweight implementation of MCAPI, that supports a much wider range of architectures than its specification normally considers. Overall, this unique combination enhances programmability by abstracting on-chip communication whilst also exposing critical parts of the target architecture to the programming language.

Local memory reservation in a predictable real-time operating systems (York)
This work concerns a simple RTOS named Carousel OS, which manages local scratchpad memory as part of the context, such that the execution time of each task is unaffected by preemption. The approach saves only the parts of the scratchpad memory which are actually required by each task, and considers the save/restore process as part of the WCET of each task, enabling a per-task tradeoff between space and WCET [YORK_2].

QoS-based resource management for safe budget overruns (UC3M, UPM)

A priority assignment protocol has been developed that allows continuous tasks to incur in safe budget overruns as a means to increase the overall quality in multimedia applications. The method is based on the contract model and budget scheduling developed by the same authors. The priority assignment is based on dual priority bands with a simple recalculation method. It is easily implementable as shown in its validation where a HOLA-QoS harness that integrates the implemented priority assignment algorithm [UC3M _1 – UC3M_12].

Component-based operating systems (Cantabria)
A methodology has been developed that allows designing hard real-time component-based applications that are going to be installed in open platforms in which the rest of the workload executed on the platform is unknown [UC_1 – UC_4]. Modelling real-time applications based on resource reservations has been proposed using two new classes named Virtual_Schedulable_Resource and Virtual_Communication_Channel. Composition mechanisms have also been developed to combine different fixed priority and EDF response-time analysis techniques for checking the schedulability of heterogeneous systems. Additionally, priority and scheduling deadline assignment techniques were combined into a new algorithm called HOSPA (Heuristic Optimized Scheduling Parameters Assignment), for optimizing the assignment of priorities and scheduling deadlines to tasks and messages in heterogeneous distributed hard real-time systems [UC_5 – UC_7].

-- The above is new material, not present in the Y3 deliverable --

3.2 Individual Publications Resulting from these Achievements

Pisa

1. [PISA_1] Giorgio Buttazzo, Enrico Bini, and Yifan Wu, "Partitioning Parallel Applications on Multiprocessor Reservations", IEEE Transactions on Industrial Informatics, Vol. 7, No. 2, pp. 302-315, May 2011.
2. [PISA_2] Mario Bambagini, Mauro Marinoni, Francesco Prosperi, and Giorgio Buttazzo, "Energy Management for Tiny Real-Time Kernels", Proceedings of the 2nd International Conference on Energy-Aware Computing (ICEAC 2011), Istanbul, Turkey, Nov. 30 – Dec. 2, 2011.
3. [PISA_3] Giorgio Buttazzo, Mauro Marinoni, and Gianluca Franchino, "Real-Time and Energy Issues in Mobile Health Monitoring Systems" Proceedings of AICA 2011, Torino, November 15-17, 2011.
4. [PISA_4] Mauro Marinoni, Mario Bambagini, Francesco Prosperi, Francesco Esposito, Gianluca Franchino, Luca Santinelli and Giorgio Buttazzo, "Platform-aware Bandwidth-oriented Energy Management Algorithm for Real-Time Embedded Systems", Proceedings of the 16th IEEE International Conference on Emerging Technology and Factory Automation (ETFA 2011), Toulouse, France, September 5-9, 2011.
5. [PISA_5] T. Cucinotta, L. Abeni, L. Palopoli, G. Lipari, "A robust mechanism for adaptive scheduling of multimedia applications," ACM Transactions on Embedded Computing Systems, Vol. 10, No. 4, November 2011.
6. [PISA_6] Giuseppe Lipari, Enrico Bini, “On the Problem of Allocating Multicore Resources to Real-Time Task Pipelines”, 4th Workshop on Compositional Theory and Technology for RealTime Embedded Systems, Vienna, Austria, November 2011.

7. [PISA_7] J. Lelli, G. Lipari, D. Faggioli, T. Cucinotta, “An efficient and scalable implementation of global EDF in Linux,” in Proc. of the 7th International Workshop on Operating Systems Platforms for Embedded Real-Time Applications (OSPERT 2011), Porto, Portugal, July 2011.
8. [PISA_8] S. Kumar, T. Cucinotta, G. Lipari, "A Latency Simulator for Many-core Systems," in Proceedings of the 44th Annual Simulation Symposium (ANSS 2011), part of the Spring Simulation Multiconference (SpringSim'11), Boston, USA, April 2011.
9. [PISA_9] T. Cucinotta and D. Faggioli, "Handling Timing Constraints Violations in Soft Real-Time Applications as Exceptions," to appear on Elsevier Journal of Systems and Software (JSS).
10. [PISA_10] T. Cucinotta, F. Checconi, L. Abeni, L. Palopoli, "Adaptive Real-Time Scheduling for Legacy Multimedia Applications," to appear on ACM Transactions on Embedded Computing Systems, Special Section on Embedded Systems for Real-Time Multimedia.
11. [PISA_11] T. Cucinotta, Fabio Checconi, George Kousiouris, Kleopatra Konstanteli, Spyridon Gogouvitis, Dimosthenis Kyriazis, Theodora Varvarigou, Alessandro Mazzetti, Zlatko Zlatev, Juri Papay, Michael Boniface, Sören Berger, Dominik Lamp, Thomas Voith, Manuel Stein. "Virtualised e-Learning on the IRMOS Real-time Cloud," to appear on Springer Service Oriented Computing and Applications.
12. [PISA_12] G. Kousiouris, T. Cucinotta, T. Varvarigou. "The Effects of Scheduling, Workload Type and Consolidation Scenarios on Virtual Machine Performance and their Prediction through Optimized Artificial Neural Networks," Journal of Systems & Software, Elsevier, Vol. 84, Issue 8, pp. 1270-1291, August 2011.
13. [PISA_13] T. Cucinotta, G. Anastasi. "A Heuristic for Optimum Allocation of Real-Time Service Workflows," (to appear) in Proceedings of the IEEE International Conference on Service-Oriented Computing and Applications (SOCA 2011), Irvine, CA, December 2011.
14. [PISA_14] T. Cucinotta, F. Checconi, D. Giani, “Improving Responsiveness for Virtualized Networking Under Intensive Computing Workloads,” in Proc. of the 13th Real-Time Linux Workshop (RTLWS 2011), Prague, Czech Republic, October 2011.
15. [PISA_15] T. Cucinotta. "Optimum Scalability Point for Parallelisable Real-Time Components," in Proceedings of the International Workshop on Synthesis and Optimization Methods for Real-time and Embedded Systems (SOMRES 2011), co-located with the 32nd IEEE Real-Time Systems Symposium (RTSS 2011), Vienna, Austria, November 2011.
16. [PISA_16] D. Kyriazis, A. Menychtas, G. Kousiouris, K. Oberle, T. Voith, M. Boniface, E. Oliveros, T. Cucinotta, S. Berger, "A Real-time Service Oriented Infrastructure," GSTF International Journal on Computing, Vol. 1, No. 2, February 2011.

Polytechnic Institute of Porto (IPP)

17. [IPP_1] José Marinho, Stefan M. Petters, “Job Phasing Aware Preemption Deferral”, In Proceedings of the Embedded Ubiquitous Computing, Melbourne, Australia, Oktober 2011
18. [IPP_2] José Marinho, Vincent Nelis, Stefan M. Petters, “Preemption Delay Analysis for Floating Non-Preemptive Region Scheduling”, In Proceedings of the in Design and Automation Conference 2012, Dresden Germany, March 2012
19. [IPP_3] Gurulingesh Raravi, Björn Andersson, Konstantinos Bletsas, “Intra-Type Migrative Scheduling of Implicit-Deadline Sporadic Tasks on Two-Type Heterogeneous

Multiprocessor", In Proceedings of the 10th Workshop on Models and Algorithms for Planning and Scheduling Problems.

20. [IPP_4] Gurulingesh Raravi, Björn Andersson, Konstantinos Bletsas, "Two-type Heterogeneous Multiprocessor Scheduling: Is there a Phase Transition?" In Proceedings of RTSOPS 2011: 2nd International Real-Time Scheduling Open Problems Seminar, Porto, Portugal, July 5, 2011.
21. [IPP_5] Gurulingesh Raravi, Björn Andersson, Konstantinos Bletsas, "Provably Good Scheduling of Sporadic Tasks with Resource Sharing on a Two-type Heterogeneous Multiprocessor Platform" In Proceedings of 15th International Conference On Principles Of Distributed Systems (OPODIS'11), December 13-16, 2011, Toulouse, France.
22. [IPP_6] Konstantinos Bletsas, Björn Andersson, "Preemption-light multiprocessor scheduling of sporadic tasks with high utilisation bound", In journal of Real-Time Systems, Volume 47 Issue 4, July 2011.

Cantabria

23. [UC_1] Laura Barros, Patricia López Martínez, and José María Drake "Design of real-time component-based applications on open platforms". 37th Euromicro Conference on Software Engineering and Advanced Applications (SEAA), Oulu (Finland), pp. 65-72, August 2011.
24. [UC_2] Laura Barros, César Cuevas, Patricia López Martínez, José María Drake, and Michael González Harbour. "Modelling real-time applications based on resource reservations". 2nd International Workshop on Analysis Tools and Methodologies for Embedded and Real-time Systems (WATERS 2011), Porto (Portugal), July 2011.
25. [UC_3] Juan M. Rivas, J. Javier Gutiérrez, J. Carlos Palencia, and Michael González Harbour. "Schedulability Analysis and Optimization of Heterogeneous EDF and FP Distributed Real-Time Systems". 23th Euromicro Conference on Real-Time Systems, Porto (Portugal), July 2011.
26. [UC_4] Michael González Harbour, J. Javier Gutiérrez, J. María Drake, Patricia López, y J. Carlos Palencia. "Modeling Real-Time Networks with MAST2". 2nd International Workshop on Analysis Tools and Methodologies for Embedded and Real-time Systems, WATERS 2011, Porto (Portugal), July, 2011, pp. 51,56.
27. [UC_5] J. Javier Gutiérrez, J. Carlos Palencia, y Michael González Harbour. "Response time analysis in AFDX networks". XIV Jornadas de Tiempo Real, JTR 2011, Madrid (Spain), February, 2011. Available at http://polaris.dit.upm.es/~str/jtr11/papers/011.pdf
28. [UC_6] Héctor Pérez Tijero, J. Javier Gutiérrez y Michael González Harbour. "Adapting the end-to-end flow model for distributed Ada to the Ravenscar profile". 15th International Real-Time Ada Workshop (IRTAW), Liébana (Spain), September, 2011
29. [UC_7] Héctor Pérez Tijero, y J. Javier Gutiérrez. "On the schedulability of a data-centric real-time distribution middleware". In press, Computer Standards and Interfaces Journal, Elsevier, August, 2011, ISSN 0920-5489.

UC3-Madrid (affiliated to Cantabria)

30. [UC3M_1] Marisol García-Valls, F. Gómez Molinero. Real-Time Reconfiguration: A Vision and its Reality . IEEE International Conference on Industrial Informatics (IEEE INDIN 2011). IEEE Computer Society Press. Caparica, Portugal. July 26 - 29, 2011.

31. [UC3M_2] M. García-Valls, P. Basanta-Val , I. Estevez-Ayres. Supporting service composition and real-time execution through characterization of QoS properties. In 6th International Symposium on Software Engineering for Adaptive and Self-Managing Systems Sponsored by ACM SIGSOFT, IEEE TCSE May 2011.
32. [UC3M_3] I. Rodríguez-López, M. García-Valls, Marisol. Architecting a Common Bridge Abstraction over Different Middleware Paradigms. 16th International Conference on Reliable Software Technologies - Ada Europe 2011. Edinburgh, UK. June, 2011.
33. [UC3M_4] M. García-Valls, P. Basanta-Val, I. Estévez-Ayres. Real-time reconfiguration in multimedia embedded systems. IEEE Transactions on Consumer Electronics, 57(3):1280-1287, August 2011
34. [UC3M_5] Marisol García-Valls, Pablo Basanta-Val, and Iria Estévez-Ayres. Real-Time Software Framework for Supporting Reconfiguration in Consumer Electronics. 11th IEEE Conference on Consumer Electronics, pp.865-866, 9-12 Jan. 2011
35. [UC3M_6] Iria Estévez-Ayres, Marisol García-Valls, Pablo Basanta-Val, Jorge Díez-Sánchez: A hybrid approach for selecting service-based real-time composition algorithms in heterogeneous environments. Concurrency and Computation: Practice and Experience 23(15): 1816-1851 (2011)
36. [UC3M_7] Iria Estévez-Ayres, Marisol García-Valls, Pablo Basanta-Val and Iván Fernández-Pacheco. Using Android Smartphones in a Service-Oriented Video Surveillance System. 11th IEEE Conference on Consumer Electronics , pp.887-888, 9-12 Jan. 2011
37. [UC3M_8] P. Basanta-Val, M. García-Valls, I. Estévez-Ayres. A dual programing model for distributed real-time Java. IEEE Transactions on Industrial Informatics, Nov. 2011.
38. [UC3M_9] P. Basanta-Val, M. García-Valls, I. Estévez-Ayres. Non-Functional Information Transmission Patterns for Distributed Real-Time Java . Software Practice and Experience. 41(12): 1409–1435 (2011)
39. [UC3M_10] Pablo Basanta-Val, Marisol García-Valls and Iria Estévez-Ayres. Fine tuning of the multiplexing facilities of Java's Remote Method Invocation. Concurrency and Computation Practice and Experience. 23(11): 236-1260, August 2011.
40. [UC3M_11] Pablo Basanta-Val, Marisol García-Valls and Iria Estévez-Ayres. Extending the Concurrency Model of the Real-Time Specification for Java. Concurrency and Computation: Practice and Experience, 23(14): 1623-1645 (2011)
41. [UC3M_12] P. Basanta-Val, M. García-Valls and I. Estévez-Ayres. Real-Time Distribution Support For Residential Gateways Based on OSGi. 11th IEEE Conference on Consumer Electronics, 747-748, 9-12 Jan. 2011

Universtity of York

42. [YORK_1] Ian Gray and Neil Audsley, "Targeting complex embedded architectures by combining the Multicore Communications API (MCAPI) with Compile-Time Virtualisation", ACM SIGPLAN/SIGBED Conference on Languages, Compilers, Tools and Theory for Embedded Systems, 2011.
43. [YORK_2] Jack Whitham & Neil Audsley, "Explicit Reservation of Local Memory in a Predictable, Preemptive Multitasking Real-time System". Proceedings of the IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS 2012), Beijing, China, 2012.

-- The above are new references, not present in the Y3 deliverable --

3.3 Interaction and Building Excellence between Partners

- **UC3M and Pisa** have started a collaboration on QoS-based resource management for multimedia applications. Currently, a PhD student from Pisa is on a one-year stay at UC3M site to collaborate on research on the integration of QoS resource managers in real-time operating systems.
- **UC3M and UPM** are collaborating in defining OS extensions in the form of QoS resource managers for multimedia systems in a distributed environment.
- **Pisa and Catalonia** continued to collaborate on event-driven scheduling to improve the efficiency and the performance of real-time controllers.
- **Pisa and Catalonia** have been working on the design of an embedded control systems laboratory experiment in order to provide a first step toward a multidisciplinary education for embedded systems engineers. By joining the know-how of SSSUP on real-time kernels and UPC on control systems, the experiment permits to flexibly work real-time systems and control systems in an integrated manner. Based on this laboratory experiment, the participation of UPC in the "ARTIST Graduate School on RT Kernels for Microcontrollers" and "PhD course on Real-Time Control Systems", reported below, have shown that this approach benefits students with different backgrounds.
- **Pisa and UPC** continued to collaborate on the SimpleRTK kernel developed at UPC [GVM10]. The kernel has been assessed by Pisa and the distribution of this kernel together with the ERIKA kernel developed by Evidence (Pisa, Italy) is under study.

-- The above is new material, not present in the Y3 deliverable --

3.4 Joint Publications Resulting from these Achievements

Pisa, TUKL, Lund, Ericsson, Evidence

1. Enrico Bini, Giorgio Buttazzo, Johan Eker, Stefan Schorr, Raphael Guerra, Gerhard Fohler, Karl-Erik Arzen, Vanessa Romero Segovia, Claudio Scordino, "Resource Management on Multicore Systems: The ACTORS Approach", IEEE Micro, Vol. 31, No. 3, pp. 72-81, May-June 2011.

Pisa & CTU-Prague

2. M. Sojka, P. Pisa, D. Faggioli, T. Cucinotta, F. Checconi, Z. Hanzalek, G. Lipari, "Modular Software Architecture for Flexible Reservation Mechanisms on Heterogeneous Resources," Elsevier Journal of Systems Architecture (JSA), Vol. 57, Issue 4, pp. 366–382, April 2011.

Pisa & UC-Berkeley

3. Andrea Sindico, Marco Di Natale, Gianpiero Panci: “Integrating SysML with Simulink using Open-source Model Transformations”. SIMULTECH 2011: 45-56.
4. Andrea Sindico, Marco Di Natale, and Gianpiero Panci: “Integrating SysML with Simulink using Open-source Model Transformations”. Proceedings of the 1st International Conference on Simulation and Modeling Methodologies, Technologies and Applications (SIMULTECH 2011), Noordwijkerhout, The Netherlands, July 29 - 31, 2011.
5. Haibo Zeng and Marco Di Natale, “Mechanisms for Guaranteeing Data Consistency and Flow Preservation in AUTOSAR Software on Multi-Core Platforms”, Proceedings of the 6th IEEE International Symposium on Industrial Embedded Systems (SIES 2011), Vasteras, Sweden, June 15-17, 2011.
6. Haibo Zeng and Marco Di Natale, “Efficient Implementation of AUTOSAR Components with Minimal Memory Usage”, Proc. of the Workshop on Synthesis and Optimization Methods for Real-time Embedded Systems (SOMRES 2011), Vienna, Austria, November 30, 2011.

Pisa & UC3-Madrid

7. G. F. Anastasi, T. Cucinotta, G. Lipari, M. Garcia-Valls, "A QoS Registry for Adaptive Real-Time Service-Oriented Applications," (to appear) in Proceedings of the IEEE International Workshop on Real-Time Service-Oriented Architecture and Applications (RTSOAA 2011), December 12-14 2011, Irvine, CA.
8. T. Cucinotta and S. V. Gogouvitis. "Real-Time Attributes in Operating Systems," Achieving Real-Time in Distributed Computing: From Grids to Clouds. IGI Global, July 2011. DOI: 10.4018/978-1-60960-827-9.
9. G. Katsaros and T. Cucinotta. "Programming Interfaces for Realtime and Cloud-based Computing," Achieving Real-Time in Distributed Computing: From Grids to Clouds. IGI Global, July 2011. DOI: 10.4018/978-1-60960-827-9.
10. E. Oliveros, T. Cucinotta, S. C Phillips, X. Yang, T. Voith, S. Middleton. "Monitoring and Metering in the Cloud," Achieving Real-Time in Distributed Computing: From Grids to Clouds. IGI Global, July 2011. DOI: 10.4018/978-1-60960-827-9.
11. S. Narasimhamurthy, M. Muggeridge, S. Waldschmidt, F. Checconi, T. Cucinotta. "Data Storage in Cloud Based Real-Time Environments," Achieving Real-Time in Distributed Computing: From Grids to Clouds. IGI Global, July 2011.
12. S. Berger, D. Lamp, M. Stein, T. Voith, T. Cucinotta, M. Bertogna. "Execution & Resource Management in QoS-aware Virtualized Infrastructures.," Achieving Real-Time in Distributed Computing: From Grids to Clouds. IGI Global, July 2011.

TUKL & Lund

13. Stefan Schorr, Anand Kotra, Gerhard Fohler, Johan Eker, Arzen, Karl-Erik, Vanessa Romero, Adaptive Resource Management in the ACTORS Framework - A live DVB-T/webcam Demo, RTSS@Work Demo Session of the IEEE Real-Time Systems Symposium, November 2011.
14. Arzen, Karl-Erik, Romero, Vanessa, Stefan Schorr, Gerhard Fohler, Reservation-Based CPU Management for Multicore Platforms, June 2011.

15. Karl-Erik Arzen, Vanessa Romero, Stefan Schorr, Gerhard Fohler, Adaptive Resource Management Made Real, Proceedings of the 3rd Workshop on Adaptive and Reconfigurable Embedded Systems (APRES 2011), Chicago, April 2011.
16. Karl-Erik Arzen, Romero, Vanessa, Mikael Kralmark, Stefan Schorr, Anand Kotra, Gerhard Fohler, Demo: Adaptive Resource Management Made Real, Proceedings of the 3rd Workshop on Adaptive and Reconfigurable Embedded Systems (APRES 2011), Chicago, April 2011.

TUKL & EPFL-Switzerland

17. Marcel Steine, Cuong Viet Ngo, Ramon Serna Oliver, Marc Geilen, Twan Basten, G. Fohler, Jean-Dominique Decotignie, Proactive Reconfiguration of Wireless Sensor Networks, 14th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems, ACM Publisher, Miami, FL, USA, October 2011.

TUKL & UNC (International Affiliated Partner)

18. Sanjoy Baruah, G. Fohler, Certification-cognizant time-triggered scheduling of mixed-criticality systems, Proceedings of IEEE Real-time Systems Symposium 2011, December 2011.

UC3M & UPM

19. Marisol García-Valls, Alejandro Alonso, Juan Antonio de la Puente. A dual-band priority assignment algorithm for QoS resource management. Future Generation Computer Systems, Elsevier. DOI: 10.1016/j.future.2011.10.005. December 2011.

Pisa & UC3M

20. Gaetano Anastasi, T. Cucinotta, G. Lipari, M. García-Valls. A QoS registry for adaptive real-time service-oriented applications. In Proc. of the IEEE International Conference on Service-Oriented Computing and Applications, pp. 1-8. California, USA. December 2011.
21. Héctor Pérez, J. Javier Gutiérrez, Esteban Asensio, Juan Zamorano, y Juan A. de la Puente. "Model-Driven Development of High-Integrity Distributed Real-Time Systems Using the End-to-End Flow Model". 37th Euromicro Conference on Software Engineering and Advanced Applications (SEAA), Oulu (Finland), Agosto, 2011.

U-Porto & U-Aveiro

22. Rui Santos, Paulo Pedreiras, Moris Behnam, Thomas Nolte, Luis Almeida. Multi-level Hierarchical Scheduling in Ethernet Switches. EMSOFT 2011, International Conference on Embedded Software. Taipei, Taiwan, 9-14 October, 2011.
23. Rui Santos, Paulo Pedreiras, Moris Behnam, Thomas Nolte, Luis Almeida. Demonstrating an Ethernet switch enhanced with hierarchical scheduling. Demo at APRES 2011, 3rd International Workshop on Adaptive and Reconfigurable Embedded Systems. Within the CPSWEEK at Chicago, USA, 11 April, 2011.
24. Ricardo Marau, Karthik Lakshmanan, Paulo Pedreiras, Luis Almeida, Raj Rajkumar, Luis Almeida. Efficient Elastic Resource Management for Dynamic Embedded Systems. ICESS

2011, The 8th IEEE International Conference on Embedded Software and Systems, Changsha, China, November 16-18, 2011.

25. Ekain Azketa, Juan Pedro Uribe, Marga Marcos, Luis Almeida, Jose Javier Gutierrez. Permutational genetic algorithm for fixed priority scheduling of distributed real-time systems aided by network segmentation. SOMRES 2011 – International Workshop on Synthesis and Optimization Methods for Real-Time Embedded Systems, a satellite event of RTSS 2011, Vienna, Austria, November 29, 2011.
26. Ekain Azketa, Juan Pedro Uribe, Marga Marcos, Luis Almeida, Jose Javier Gutierrez. Genetic Algorithm for the Optimized Assignment of Priorities to Tasks and Messages in Distributed Real-Time Systems. ICESS 2011, The 8th IEEE International Conference on Embedded Software and Systems, Changsha, China, November 16-18, 2011.
27. 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, 58(3):1061-1071. DOI: 10.1109/TIE.2010.2049711, March 2011.
28. Javier Silvestre-Blanes, Paulo Pedreiras, Ricardo Marau. Bandwidth Mapping Algorithms in Distributed Media Control Applications. ETFA 2011, The 16th IEEE Conference on Emerging Technologies for Factory Automation, Toulouse, France, 5-9 Sep 2011.

-- The above are new references, not present in the Y3 deliverable --

3.5 Keynotes, Workshops, Tutorials

Keynotes

1. L. Almeida. A Network-centric Perspective on Cyber-Physical Systems. Dagstuhl Seminar on Cyber-Physical Systems, Dagstuhl, Germany, 2-4 November 2011.
2. L. Almeida. On Cyber-Physical Systems. EU-US Workshop on Networked Monitoring and Control, Brussels, Belgium, 28 June 2011.
3. L. Almeida. When Time Becomes Real. Seminar at the University of Aveiro, within the scope of the MAP-Tele PhD program, 18 Feb 2011.
4. Giorgio Buttazzo, “Supporting real-time applications on multicore platforms”, Keynote Talk at the 6th IEEE International Symposium on Industrial Embedded Systems (SIES 2011), Vasteras, Sweden, June 15-17, 2011.
5. Giuseppe Lipari, “Component-based analysis of real-time systems”, Keynote Talk at the 19th International Conference on Real-Time and Networked Systems (RTNS 2011), Nantes, France, September 29-30, 2011.
6. Tommaso Cucinotta, “SOOS: Issues in Large Scale Scheduling of Distributed Applications,” at the Conference on Computing Architectures, Software tools and nano-Technologies for Numerical and Embedded Scalable Systems (CASTNESS 2011), Rome, January 17-18, 2011.
7. Marco Di Natale, “From heterogeneous models to code: Challenges in the development of modern real time automotive software”, at the SAE International Vehicle Battery Summit (SAE 2011), Shanghai, China, November 14-15, 2011.

8. Marco Di Natale, “From analysis to optimization in the deployment of real-time distributed functions in modern automotive systems”, at the Workshop on Software Synthesis (WSS 2011), Taipei, Taiwan, October 14, 2011.
9. Marisol García Valls. "Real-Time Middleware". University Carlos III de Madrid. Master on Aricrat Systems Integration. EADS-UC3M. May 2011.

Participation on Conference and Workshops

1. RTAS 2011 - IEEE Real-Time and Embedded Technology and Applications Symposium, Chicago, Illinois, USA, April 12-14, 2011.
   URL: http://www.rtas.org/
2. RTSS 2011 - IEEE Real-Time Systems Symposium, Vienna (Austria), November 30th, - December 2nd 2011.
   URL: http://www.rtss.org/
3. ECRTS 2011 - Euromicro Conference on Real-Time Systems, Porto, Portugal, July 6-8, 2011.
   URL: http://ecrts11.ecrts.org
4. ETFA 2011 - IEEE International Conference on Emerging Technologies and Factory Automation, Toulouse, France, September 5-9, 2011.
   URL: http://www.etfa-2011.org/
5. RTCSA 2011 - IEEE International Conference on Embedded and Real-Time Computing Systems and Applications, Toyama, Japan, August 28-31, 2011.
   URL: http://www.jaist.ac.jp/rtcsa2011/
6. RTNS 2011 - International Conference on Real-Time and Network Systems, Nantes, France, September 29-30, 2011.
   URL: http://rtns2011.irccyn.ec-nantes.fr/
7. OSPERT 2011 – The 7th International Workshop on Operating Systems Platforms for Embedded Real-Time Applications, Porto, Portugal, July 5th, 2011.
   URL: http://www.seas.gwu.edu/~gparmer/ospert11/
8. WCET 2011 – The 11th International Workshop on Worst-Case Execution Time Analysis, Porto, Portugal, July 5th, 2011.
   URL: http://cs.furman.edu/~chealy/wcet2011/
9. RTN 2011 – The 10th International Workshop on Real-Time Networks, Porto, Portugal, July 5th, 2011.
   URL: http://www.rtn2011.org/
10. WATERS 2011 – 2nd International Workshop on Analysis Tools and Methodologies for Embedded and Real-time Systems, Porto, Portugal, July 5th, 2011.
    URL: http://retis.sssup.it/waters2011/
11. RTSOPS 2011: 2nd International Real-Time Scheduling Open Problems Seminar, Porto, Portugal, July 5th, 2011.
    URL: http://www.cs.wayne.edu/~fishern/Meetings/rtsops2011/
12. CPSWEEK 2011 – The 5th Cyber-Physical Systems Week, Chicago, Illinois, USA, April 12-14, 2011.
    URL: http://cpsweek2011.cs.illinois.edu/
13. APRES 2011: 3rd Workshop on Adaptive and Reconfigurable Embedded Systems, Chicago, Illinois, USA, April 11, 2011
    URL: http://apres2011.uwaterloo.ca/

14. ICCPS 2011: 2nd ACM/IEEE International Conference on Cyber-Physical Systems, Chicago, Illinois, USA, April 11, 2011.
    URL: http://precise.seas.upenn.edu/events/iccps11/
15. CRTS 2011 – 4th Workshop on Compositional Theory and Technology for Real-Time Embedded Systems, Vienna, Austria, November 29th, 2011.
    URL. http://www.rapitasystems.com/crts2011/
16. AVICPS 2011: 2nd International Workshop on Analytic Virtual Integration of Cyber-Physical Systems, Vienna, Austria, November 29th, 2011.
    URL: http://www.analyticintegration.org/
17. WCTT 2011: 1st International Workshop on Worst-case Traversal Time, Vienna, Austria, November 29th, 2011.
    URL: http://www.wctt.info/
18. SOMRES 2011: Workshop on Synthesis and Optimization Methods for Real-time Embedded Systems, Vienna, Austria, November 29th, 2011.
    URL: http://retis.sssup.it/synthesys/
19. RTSS@Work 2011: Real-Time SystemS at Work - Open Demo Session of Real-Time Techniques and Technologies, Vienna, Austria, November 29th, 2011.
    URL: http://webpages.cister.isep.ipp.pt/~smp/RTSS@Work/
20. DATE 2011 – Design, Automation & Test in Europe, Grenoble, France, March 14 – 18, 2011.
    URL: http://www.date-conference.com/
21. HSCC 2011: 14th ACM International Conference on Hybrid Systems: Computation and control, Chicago, Illinois, USA, April 12-14, 2011.
    URL: http://hscc2011.cs.sunysb.edu/
22. NeRES 2011: 2nd Workshop on Networks for Real-time Embedded Systems. Porto,Portugal, 10-11 November, 2011.
    URL: http://paginas.fe.up.pt/~ftt/neres2011/
23. MED 2011: The 19th Mediterranean Conference on Control and Automation, Corfu, Greece, June 20-23 2011.
    URL: http://www.med2011.org/

Tutorials

1. **Graduate Course on Combinatorial Optimization**
   Scuola Superiore Sant'Anna, Pisa, Italy – October-November 2011

   Objectives: The course covered several techniques of combinatorial optimization, namely complete methods such as Constraint Programming, Integer Linear Programming, Dynamic Programming and incomplete methods that go from simple local search to more sophisticated meta-heuristics. A final lecture on Hybrid Scheduling showed integrated methods for scheduling problems, in particular, allocation and scheduling of task graph applications on MPSoCs.

   Organizers: Giorgio Buttazzo (Scuola Superiore Sant'Anna), Michela Milano (Univ. of Bologna, Italy).

2. **ARTIST Graduate Course on Real-Time Kernels for Microcontrollers**
   Scuola Superiore Sant'Anna, Pisa, Italy – June 13-17, 2011

Objectives: The course was aimed at providing the fundamentals concepts of real-time computing systems, including scheduling, resource management and timing analysis; introducing the OSEK/VDX standards, taking as a reference implementation the Erika Enterprise kernel; showing how to apply such concepts in practice, with examples based on the Flex platform and the Microchip dsPIC DSC microcontrollers; teaching participants how to develop simple control applications using Erika Enterprise with code generation from functional models.

Organizers: Giorgio Buttazzo (Scuola Superiore Sant'Anna), Pau Marti (Technical University of Catalonia, Barcelona, Spain), Ettore Ricciardi (ISTI-CNR, Pisa).

URL: http://www.artist-embedded.org/artist/-ARTIST-Graduate-School-on-RT,1200-.html

3. Graduate Course on Android Framework
Scuola Superiore Sant'Anna, Pisa, Italy – November-December 2011

Objectives: The course explained how to develop Android systems: from application bound entities, to the innovative (pseudo-)distributed IPC model, going through the key features of this framework; understanding how quality applications for Android should be developed. The framework analysis showed how an sample feature has been designed from the Java API down to the Android driver stub, enabling the student to extend the Android framework and to export its dedicated SDK.

Organizers: Giorgio Buttazzo (Scuola Superiore Sant'Anna), Alberto Panizzo (Amarula Solutions, Italy).

4. Workshop on Real-Time System Models for Schedulability Analysis
University of Cantabria, Santander, Spain - February 7-8, 2011

Objectives: Present existing models of real-time systems, and propose extensions to fill the gaps that are required to cover state-of-the-art hardware platforms, operating systems, and scheduling techniques used in practice to develop real-time applications.

Organizers: Michael González Harbour (Universidad de Cantabria).

5. 15th International Real-Time Ada Workshop (IRTAW-15)
Liébana (Cantabria), Spain, September 2011

The 15th International Real-Time Ada Workshop (IRTAW-15) took place on September 14-16 of 2011 in Liébana (Cantabria), Spain. Since the late Eighties the International Real-Time Ada Workshop series has provided a forum for identifying issues with real-time system support in Ada and for exploring possible approaches and solutions, and has attracted participation from key members of the research, user, and implementer communities worldwide. Recent IRTAW meetings have significantly contributed to the Ada 2005 standard and to the proposals for Ada 2012, especially with respect to the tasking features, the real-time and high-integrity systems annexes, and the standardization of the Ravenscar profile. This particular meeting was organized by the University of Cantabria and received 22 participants from different countries in Europe and North America. The discussions were centred around multiprocessor real-time scheduling, multiprocessor resource control protocols, language profiles, application frameworks, and concurrency in Ada 2012. Some of the results of the workshop were submitted to the standardizaition bodies producing the Ada 2012 standard, and some others were captured in the proceedings to generate new work for future standardization phases.

URL: http://www.artist-embedded.org/artist/IRTAW-15,2204.html

Organization of Conference and Workshops

1. Giorgio Buttazzo has been General Co-Chair of the 8th IEEE International Conference on Embedded Software and Systems (ICESS 2011), Changsha, China, November 16-18, 2011.
2. Giorgio Buttazzo has been nominated General Chair for the 24th Euromicro Conference on Real-Time Systems (ECRTS 2012), Pisa, Italy, June 11-13, 2012.
3. Marco Di Natale has been nominated Program Chair of the 18th IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS 2012), Beijing China, April 16-19 2012.
4. Giuseppe Lipari has been Program Chair at the 15th International Conference on Principles of Distributed Systems (OPODIS 2011), Dec. 12-16, 2011, Toulouse, France.
5. Giuseppe Lipari has been Program Co-Chair of the 2nd International Workshop on Analysis Tools and Methodologies for Embedded and Real-time Systems (WATERS 2011), Porto, Portugal, July 5th, 2011.
6. Giuseppe Lipari has been nominated Track Chair for the track on Real-time, Networked and Dependable Systems, at the ACM Design and Automation Conference (DATE 2012), March 2012, Dresden, Germany.
7. Marco Di Natale has been Program Co-Chair at the 8th IEEE International Conference on Embedded Software and Systems (ICESS 2011), Changsha, China, November 16-18, 2011.
8. Marco Di Natale has been Program Co-Chair of the Work-In-Progress Session at the 6th IEEE International Symposium on Industrial Embedded Systems (SIES 2011), Vasteras, Sweden, June 15-17, 2011.
9. Tommaso Cucinotta has been Program Co-Chair of the Fourth IEEE International Workshop on Real-Time Service-Oriented Architecture and Applications (IEEE RTSOAA 2011), December 12th-14th, 2011, Irvine, USA
10. Marco Di Natale has been nominated Track Co-Chair for the track on Transportation and Energy, at the ACM Design and Automation Conference (DATE 2012), March 2012, Dresden, Germany.
11. Marco Di Natale has been nominated Track Co-Chair for the track on Real-Time and (Networked) Embedded Systems, at the 17th International Conference on Emerging Technologies and Factory Automation (ETFA 2012), September 17 - 21, 2012, Krakow, Poland.
12. Tommaso Cucinotta has been Program Co-Chair of the 2nd International Workshop on Analysis Tools and Methodologies for Embedded and Real-time Systems (WATERS 2011), Porto, Portugal, July 5th, 2011.
13. Enrico Bini has been Program Chair of the Work-In-Progress Session of the 23rd Euromicro Conf. on Real-Time Systems (ECRTS 2011), Porto, Portugal, July 6 - 8, 2011.

14. Enrico Bini has been Workshops Chair at the 32nd IEEE Real-Time Systems Symposium (RTSS 2011), Vienna, Austria, November 29 - December 2, 2011.

-- The above is new material, not present in the Y3 deliverable --

4. Internal Reviewers for this Deliverable

- **Karl-Erik Arzen** (University of Lund)
- **Luis Almeida** (University of Porto, Portugal)