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

Activity Progress Report for Year 2

Resource-aware Operating Systems

Cluster:

Operating Systems and Networks

Activity Leader:

Prof. Giorgio Buttazzo (Scuola Superiore S. Anna)
http://feanor.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	December 18th 2009

Table of Contents

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 Y1 deliverable --

No changes with respect to Year 1.

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: Liesbeth Steffens – NXP (The Netherlands)
Role: real-time kernels for media processing applications.

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: memory management in real-time Java middleware.

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

-- Changes wrt Y1 deliverable --

No changes with respect to Year 1.

1.3 Starting Date, and Expected Ending Date

Starting date: January 1st, 2008

Ending date. December 31st, 2010.

-- Changes wrt Y1 deliverable --

No changes with respect to Year 1.

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 Y1 deliverable --

No changes with respect to Year 1.

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 Y1 deliverable --

No changes with respect to Year 1.

1.6 Technical Description: Joint Research

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 Y1 deliverable --

Effort to include deadline scheduling in the Linux kernel, extend resource reservations to multicore platforms and investigate hierarchical scheduling for component-based design.

1.7 Work achieved in Year 1

- **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 Y1 deliverable --

This section was already presented in the Y1 deliverable, in sections 1.7 and 3.1.

1.8 Problems Tackled in Year 2

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.

-- Changes wrt Y1 deliverable --

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

2. Summary of Activity Progress in Year 2

2.1 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 tacking timing and precedence constratins into account. The method generalizes the resource reservation abstraction to multicore systems and allows allocating a fraction of the available resources to a given application.

The tool is available from:
http://retis.sssup.it/tools/mrr/index.html
http://xen1.sssup.it/yifan/mrr/index.html

Deadline Scheduling on Linux (Pisa, Evidence)
The Linux kernel has been extended to support deadline-based scheduling algorithms, like Earliest Deadline First (EDF). The contribution to the community has been twofold: On one hand we have designed and developed a new POSIX-like scheduling policy, called SCHED_DEADLINE, implemented as a novel scheduling class, like the POSIX-specified ones already present in the kernel; on the other hand we have extended the existing bandwidth control mechanism (called throttling), basing it on a CBS server.
The new scheduling policy, SCHED_DEADLINE, enables users to create EDF tasks, introducing an ad-hoc interface to fully specify task parameters. The policy supports the well-known sporadic task model, and takes priority over the old SCHED_FIFO and SCHED_RR policies. Applications need to be modified to take advantage of it, but they receive all the benefits of EDF scheduling.
The throttling reimplementation can be used to control the bandwidth assigned to legacy tasks making use of a standard grouping mechanism; tasks are still scheduled with a fixed priority algorithm within the same group (so they need no modification to be executed) but throttling, which previously was just a safety measure against runaway tasks, can now be used to enforce timing guarantees among groups.

A general description of EDF on Linux is available from:
http://retis.sssup.it/?q=node/35

SCHED_DEADLINE is available from:
http://www.evidence.eu.com/sched_deadline.html

The throttling reimplementation is available from:
http://feanor.sssup.it/~fabio/linux/edf-throttling/

Kernel support for power-aware embedded systems (Pisa)
The RETIS group of Pisa stated to extend the Erika kernel to support the development of real-time applications with energy constraints. In particular, a kernel library is being developed to help the programmer to specify resource requirements, energy and timing constraints, and select the most appropriate operational modes for the processor and the peripheral devices.

Simulation of real-time applications (Pisa, Bologna, U-Saarland, U-Dormunt)
The Erika real-time kernel (developed by Evidence) has been ported on top of the MPARM simulator (developed by the Univ. of Bologna) to evaluate the effect of different architecture solutions and resource management policies on real-time applications.
MPARM is a SystemC based, cycle-accurate simulation environment for multiprocessor system-on-chip (MPSoC) architectures. It is specifically designed to reflect embedded system characteristics for asymmetric multiprocessing systems (i.e., general purpose, DSP, VLIW can coexist). The simulator is intended as a tool to explore different architectural choices since early design phases and can measure the performance of different architecture components, such as bus usage, cache and processor performance, power requirements, etc.
Erika is a real-time kernel compliant with the OSEK/VDX standard, designed to have minimal memory requirements. It supports mono and multi processor systems and is available for several small 8 to 32 bit microcontrollers and platforms from Altera's NIOS II Softcores to the dsPIC DSC micro-controllers. It supports preemptive and non-preemptive execution, shared resources, fixed priorities and deadline scheduling.
In particular, MPARM was configured to simulate an ARM7 core (the ARM7TDMI processor Samsung KS32C50100 in the Evaluator-7T board). Porting ERIKA to MPARM required the modification of the context switching code to support the processor, the Startup, and interrupt controller configuration, the irq handling code, and the API for useful simulation support functions.
ERIKA: http://erika.sssup.it/
MPARM: http://www-micrel.deis.unibo.it/sitonew/research/mparm.html

Execution of multiple real-time operating systems on a multicore architecture (Pavia)
This research activity focused on the development of a software computing technology for allowing the concurrent execution of different real-time operating systems on a multicore architecture, where each core is assigned a different running kernel. The software has been implemented using an open source codebase, namely the Linux kernel and the S.Ha.R.K. real-time kernel.

Distributed real-time control of a robotic manipulator (Pavia)
The team at the University of Pavia is working at applying real-time computing techinques to the distributed control of a robotic manipulator. The goal is to remotely control the manipulator over a dedicated TCP/IP connection, in order to move in a working environment with obstacles. The real-time control application is developed using the S.Ha.R.K. real-time kernel. It requires to be remotely interfaced with different kind of sensors and actuators, such as DC motors, resolvers, force sensors, and cameras. The enforcement of strict timing constraints on the computing tasks is mandatory to achieve adequate performance of the robotic application.

Real-time applications on the educationl kit (Pisa, Pavia)
The RETIS Lab of Pisa and the Robotic Lab of the University of Pavia collaborated in the development of different real-time applications on the FLEX-ERIKA educational control kit provided by Evidence. Applications include: the control of a ball-and-plate device, using infrared sensors as sensing devices; coordination of a team of robot vehicles; wireless broadcasting of audio signals; real-time processing of audio signals; the development of an Inertial Measurement Unit for the navigation of autonomous robots; the development of a segway; the use of a force sensor for the automatic motion of a 6 degree-of-freedom robotic manipulator; a localization technique based on real-time sensor fusion between inertial measurements (acceleromenters) and wireless signal strength information; and real-time monitoring of power consumption in buildings.
The FLEX board: http://www.evidence.eu.com/content/view/125/206/
The ERIKA Enterprise: http://www.evidence.eu.com/content/view/27/254/
Applications: http://robot.unipv.it/index.php/scientific-activities/research-topics

RTNS simulation package (Pisa).
Simulation can play a key role, together with analytical models, for validating a system against the QoS it must guarantee. Pisa is continuing to work on RTNS a publicly available free tool to simulate Operating System (OS) aspects in wireless distributed applications. The tool extends the well-known NS-2 simulator with models of the CPU, RealTime OS and application tasks, to take into account delays introduced by computation and I/O from peripherals. This package can be used to efficiently co-design the kernel and network profiles before deployment, especially when the number of nodes is such that analytical analysis is inapplicable.

URL: http://rtns.sssup.it

The IEEE 802.15.4 network stack in ERIKA Enterprise (Pisa, Porto).
The philosophy underlying the standardization of the IEEE 802.15.4 protocol for Low Rate-Wireless Personal Area Network is to support the large majority of present, planned, and envisioned networked applications deployed through cost-effective and autonomous wireless nodes. The standard appropriately encompasses features coming from scheduled-based (through the so-called Guaranteed Time Slots, GTS) and contention-based (through the CSMA/CA mechanism) medium access paradigms thus supporting real-time and best-effort traffic types. Moreover particular attention is devoted to energy policy issues (for battery life extension) allowing the nodes in a network to synchronously sleep and wake up according to a broadcasted schedule.

Erika + OpenZB software suite: http://www.evidence.eu.com/content/view/129/380/

Video decoding on multicore platforms (TUKL).
TUKL continued to work on strategies for distributing functional parts of video decoding on multicore platforms with particular focus on synchronization issues. Previous work focused on MPEG-2, which has now been extended to MPEG4.

Operating system abstractions for wireless sensor networks (TUKL, Philips).
TUKL and Philips have been working on an operating system abstraction layer for portable applications in wireless sensor networks. An API has been defined and ports to FreeRTOS and MantisOS implemented.

Execution time monitoring and interrupt handlers (Cantabria)
In hard real-time systems it is essential to monitor the execution times of all tasks and detect situations in which the estimated worst-case execution time (WCET) is exceeded. The accuracy of the POSIX execution time clocks and timers can be enhanced significantly if they are not arbitrarily charged with the execution times of interrupt service handlers, as occurs in simple implementations. We have implemented support for measuring the execution time of interrupt handlers separately from the execution time of user threads. We have measured the overheads incurred by the implementation of the new services, and found that they are very low in comparison with the execution times of common interrupt handlers.

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

2.2 Individual Publications Resulting from these Achievements

Pisa

1. Fabio Checconi, Tommaso Cucinotta, Dario Faggioli, Giuseppe Lipari, "Hierarchical Multiprocessor CPU Reservations for the Linux Kernel", in Proceedings of the 5th International Workshop on Operating Systems Platforms for Embedded Real-Time Applications (OSPERT 2009), Dublin, Ireland, June 2009.
2. Tommaso Cucinotta, Dario Faggioli, Alessandro Evangelista, "Exception-Based Management of Timing Constraints Violations for Soft Real-Time Applications", in Proceedings of the 5th International Workshop on Operating Systems Platforms for Embedded Real-Time Applications (OSPERT 2009), Dublin, Ireland, June 2009.
3. Tommaso Cucinotta, Gaetano Anastasi, Luca Abeni "Respecting temporal constraints in virtualised services", in Proceedings of the 2nd IEEE International Workshop on Real-Time Service-Oriented Architecture and Applications (RTSOAA 2009), Seattle, Washington, July 2009.
4. Dario Faggioli, Fabio Checconi, Michael Trimarchi, Claudio Scordino, “An EDF scheduling class for the Linux kernel”, 11th Real-Time Linux Workshop (RTLW), Dresden, Germany, September 2009.

Pavia (affiliated to Pisa)

1. 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.
2. 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.
3. 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. ISBN: 978-88-6055-415-4.
4. 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. ISBN: 978-88-6055-415-4.

York

1. K. Yu and N. Audsley, “A Generic and Accurate RTOS-centric Embedded System Modelling and Simulation Framework”, Fifth UK Embedded Forum (UKEF09), Sept, 2009.
2. J. Whitham and N. Audssley, “Implementing Time-predictable Load and Store Operations”, Proc. EMSOFT, 265-274, 2009.

Cantabria

1. 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.

TUKL

1. Ramon Serna Oliver, Ivan Shcherbakov, Gerhard Fohler, An Operating System Abstraction Layer for Portable Applications in Wireless Sensor NetworksProceedings of The 25th ACM Symposium on Applied Computing (SAC 2010), ACM, Sierre, Switzerland, March 2010.

UC3M (affiliated Cantabria)

1. 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.
2. 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. Acceptance rate: 20%.
3. 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 Service Oriented Architectures in Converging Networked Environments. (IEEE SOCNE'09) Bradford, UK. May 2009.
4. P. Basanta-Val, M. García-Valls, I. Estévez Ayres. Simple asynchronous remote invocations for distributed real-time Java. To appear in IEEE Trans. On Industrial Informatics, August 2009. Impact factor: 2,3 (Q1).
5. P. Basanta-Val, M. García-Valls, I. Estévez Ayres. No heap remote objects for distributed real-time Java. ACM Trans. on Embedded Computing Systems.
6. 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.
7. I. Estévez-Ayres, M. García-Valls, and P. Basanta. Towards Distributed Composition of Real-Time Service-Based Applications . In Proc. of the 4th International IEEE Service Oriented Architectures in Converging Networked Environments. (IEEE SOCNE'09) Bradford, UK. May 2009.

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

2.3 Interaction and Building Excellence between Partners

- Pisa and Evidence have collaborated to extend the Linux kernel to support deadline-based scheduling and resource reservations. They developed a new POSIX-like scheduling class, called SCHED_DEADLINE and extended the existing bandwidth control mechanism (called throttling), using resource reservations.
- Pisa, TUKL and Lund addressed the problem of executing real-time applications on top of multicore platforms, extending the resource reservation framework from uniprocessor to multicore systems. In particular, Pisa focused on partitioning and allocation algorithms for optimizing the consumed bandwidth, Lund investigated feedback scheduling techniques to

achieve adaptivity for dynamic applications with highly variable bandwidth requirements, and TUKL specified the architecture framework and the various interfaces for adaptive resource management.

- Pisa and Saarland University started a joint research for investigating the effects of preemptive scheduling on worst-case computation times. In particular, Pisa developed a limited preemptive approach to take advantage of preemptions without paying extra blocking delays due to non preemptive regions. Exact schedulability analysis also allows computing the longest non-preemptive region for each task that maintain the task set feasibile.
- Pisa and Bologna worked together to port the Erika real-time kernel (developed by Evidence) on top of the MPARM architecture emulator (developed by the Univ. of Bologna) to evaluate the effects of different architecture solutions and resource management policies on tasks execution times. Extensive simulation experiments are being performed to understand how the cache behaviour is affected by scheduling algorithms.
- Pisa and Catalonia (affiliated to Lund) are continuing to collaborate on event-driven scheduling. The objective of this research is to exploit the knowledge of the system to be controlled to relax the periodicity contraints of control activities to reduce resource utilization.
- Pisa and Catalonia (affiliated to Lund) are also collaborating for developing educational control applications for disseminating real-time control methodologies using practical laboratory experiences.
- Pisa and Porto are collaborating to implement the IEEE 802.15.4 standard as a communication module of ERIKA Enterprise to enable node-to-node real-time communication in WSN scenarios. This work has been pushed by the research community that failed in such an effort without the support of a real-time kernel. The preliminary tests on node performances proved that ERIKA-based networked applications outperform the corresponding codes implemented on top of non real time kernels like TinyOS.
- The Research Unit at the University of Pavia has active international collaborations with the University of Porto (Portugal, Prof. Luis Almeida), University of Malardalen (Sweden, Prof. Thomas Nolte) and University of Aveiro (Portugal, Prof. Nuno Lau).
- Pisa, Evidence, Embedded Solutions, and Microchip Technology, started a collaboration to produce a new educational control kit consisting of a balancing ball-and-plate device, using a touch screen as a sensor for detecting the ball position.

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

2.4 Joint Publications Resulting from these Achievements

1. Luca Abeni, Luigi Paolopoli, Claudio Scordino, Giuseppe Lipari, “Resource Reservations for General Purpose Applications”, IEEE Transactions on Industrial Informatics, Volume 5:1, Pages 12-21, February 2009.
2. Konstanteli, K., Kyriazis, D., Varvarigou, T., Cucinotta, T. and Anastasi, G. "Real-Time Guarantees in Flexible Advance Reservations". In Computer Software and Applications Conference, 2009. COMPSAC '09. 33rd Annual IEEE International, pages 67-72, 2009.
3. Cucinotta, T., Anastasi, G. and Abeni, L. "Respecting Temporal Constraints in Virtualised Services". In Computer Software and Applications Conference, 2009. COMPSAC '09. 33rd Annual IEEE International, pages 73-78, 2009.

4. Tommaso Cucinotta, Luigi Palopoli, "QoS Control for Pipelines of Tasks Using Multiple Resources," IEEE Transactions on Computers, IEEE Computer Society Digital Library, 23 July 2009.
5. Tommaso Cucinotta, Antonio Mancina, Gaetano Anastasi, Giuseppe Lipari, Leonardo Mangeruca, Roberto Checcozzo, Fulvio Rusinà, "A Real-time Service-Oriented Architecture for Industrial Automation," IEEE Transactions on Industrial Informatics, Vol. 5, n. 3, pp. 267-277, August 2009.
6. Tommaso Cucinotta, Kleopatra Konstanteli, Theodora Varvarigou "Advance Reservations for Distributed Real-TimeWorkflows with Probabilistic Service Guarantees", to appear in Proceedings of the IEEE International Conference on Service-Oriented Computing and Applications (SOCA 2009), Taipei, Taiwan, December 2009.
7. Tommaso Cucinotta, Luca Abeni, Luigi Palopoli, Fabio Checconi "The Wizard of OS: a Heartbeat for Legacy Multimedia Applications", in Proceedings of the 7th IEEE Workshop on Embedded Systems for Real-Time Multimedia, Grenoble (ESTImedia 2009), Grenoble, October 2009.
8. Tommaso Cucinotta, Giuseppe Lipari, Luigi Palopoli, Luca Abeni, Rodrigo Santos "Multi-level feedback control for Quality of Service Management", in Proceedings of the 14th IEEE International Conference on Emerging Technologies and Factory Automation, Palma de Mallorca (ETFA 2009), Spain, September 2009.
9. Fabio Checconi, Tommaso Cucinotta, Manuel Stein "Real-Time Issues in Live Migration of Virtual Machines", in Proceedings of the 4th Workshop on Virtualization and High-Performance Cloud Computing (VHPC 2009), Delft, The Netherlands, August 2009.
10. Kleopatra Kostanteli, Dimosthenis Kyriazis, Theodora Varvarigou, Tommaso Cucinotta, Gaetano Anastasi, "Real-time guarantees in flexible advance reservations", in Proceedings of the 2nd IEEE International Workshop on Real-Time Service-Oriented Architecture and Applications (RTSOAA 2009), Seattle, Washington, July 2009.
11. Luigi Palopoli, Tommaso Cucinotta, Luca Marzario, Giuseppe Lipari, "AQuoSA - adaptive quality of service architecture." Softw., Pract. Exper. 39(1), pp. 1-31, 2009.
12. A. Camacho, P. Martí, M. Velasco and E. Bini, “Implementation of Self-triggered Controllers”. In Demo Session of 15th IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS09), San Francisco, CA, USA, April 2009.
13. 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. Impact factor: 1,9 (Q1).
14. 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. Impact factor: 2,3 (Q1).
15. 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 Techonologies and Factory Automation (ETFA), Palma De Mallorca, Spain, 22-26, September, 2009.
16. 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.
17. Gianluca Franchino, Giorgio Buttazzo, and Tullio Facchinetti, "Token Passing Techniques for Hard Real-Time Communication", in Factory Automation, ISBN 978-953-7619-42-8, to appear.

2.5 Keynotes, Workshops, Tutorials

Keynote: Enrico Bini (SSSA): "Real-Time Scheduling for Control Systems"
17th International Conference on Real-Time and Network Systems (RTNS 2009)
Paris, France – October 26th, 2009

The talk presented an overview of the techniques that can be used to design control systems taking performance requirements and schedulability constraints into account.

http://rtns09.ece.fr/index_fichiers/Schedule.html - Speaker

Tutorial: Graduate Course on Embedded Control Systems: Theory and Practice
Scuola Superiore Sant'Anna, Pisa, Italy – June 8-12, 2009

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), Karl-Erik Arzen (Lund University), Luis Almeida (University of Porto), Ettore Ricciardi (ISTI-CNR, Pisa).

URL: http://www.artist-embedded.org/artist/ARTIST-Embedded-Control-2009.html

Workshop: CyberRescue@RTSS2009
Washington DC, December 1, 2009

Objectives: CyberRescue@RTSS2009 is an international competition for simulated robotic agents that is held in conjunction with the Real-time Systems Symposium (RTSS'09) in Washington DC. The event is sponsored by the ArtistDesign NoE.

Organizers: Tullio Facchinetti (University of Pavia) and Luis Almeida (University of Porto).

URL: http://robot.unipv.it/cyberrescue-RTSS09/

Workshop: XII Spanish workshop on Real-Time Systems
UC3-Madrid, 5-6 february 2009
Yearly workshop for researchers on Real-Time Systems

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

PROMETEO (Spanish National Industry Platform on Embedded Systems)
Madrid, SPAIN – October 28th, 2009
Presentation of iLAND project to National researchers and Industry.

ARTEMIS Autum Event (Held joinly with ITEA2)
Madrid, SPAIN – October 29-30, 2009
Presentation of iLAND project to European and National Authorities and Industrial and Research community.

3. Milestones, and Future Evolution

3.1 Problem to be Tackled over the next 12 months (Jan 2010 – Dec 2010)

In the next 12 months the cluster will investigate the follwing issues:

- Continue the development of software modules to support real-time control applications on the educational kit.
- Evaluate the partitioning algorithm on a specific multicore platform, using the Erika real-time kernel and investigate the allocation of dynamic applications with highly variable resource requirements.
- Investigate a method for positioning preemption points in the task code to reduce the cache related pre-emption delays in task execution times.
- Consider the effect of shared resources and non-preemptive regions in resource reservations and evaluate their effects on temporal isolation.
- Extend the Erika kernel to support the development of real-time applications with energy constraints.
- Improve resource usage strategies inside real-time middleware platforms.

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

Current and Future Milestones

Year 1:

1. **Model resource reservation on multi-core architectures (Achieved)**. The concept of resource reservation has been extended to multi-core architectures, with the objective of achieving temporal protection among different applications.
2. **Educational kit for embedded control applications (Achieved)**. An educational kit for embedded systems, based on Microchip dsPIC technology, has been developed at the Scuola Superiore Sant'Anna of Pisa, in collaboration with Pavia, Evidence, Embedded Solutions, and Microchip Technology. It consists of a number of modules that can easily be composed depending on specific application purposes. The modules include a set of libraries to simplify the access to the hardware devices (sensors, servomotors, wireless modules) and a number of sample real-time control applications that can be easily replicated by the users.
3. **XtratuM hypervisor (Achieved)**. The XtratuM hypervisor has been developed at the Technical University of Valencia and has substituted the RTLinux distribution to provide full temporal and spatial isolation to real-time partitions and Linux partition. In particular, XtratuM has been ported to Leon2 processors (the processor used by ESA in the space) under a project funded by CNES (France) following the ARINC 653 standard. XtratuM can now be used as a test platform for the techniques proposed in this activity.

Year 2:

1. **Partitioning real-time applications on multi-core platforms (Achieved)**. The problem of partitioning parallel real-time applications on top of multicore platforms has been addressed for exploiting the available parallelism offered by modern architectures and optimizing the required resources. A method has also been provided for guaranteeing the timing constraints of the application.
2. **Deadline Scheduling on Linux (Achieved)**. The Linux kernel has been extended to support deadline-based scheduling algorithms and resource reservations.
3. **Software support for embedded control applications (Achieved)**. A set of modules to simplify the development of real-time control applications on the educational control kit has been developed at the Scuola Superiore Sant'Anna of Pisa, in collaboration with Pavia, Evidence, Embedded Solutions, and Microchip Technology. The modules include a set of libraries to simplify the access to the hardware devices (sensors, servomotors, wireless modules), a number of sample real-time control applications that can be easily replicated by the users, a real-time communication MAC protocol for wireless sensor networks.

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

3.2 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:

FRESCOR - Framework for Real-time Embedded Systems based on COntRacts

FRESCOR is a consortium research project funded in part by the European Union's Sixth Framework Programme (FP6/2005/IST/5-034026). The following ArtistDesign partners are involved: University of Cantabria, University of York, Scuola Superiore Sant'Anna, Technical University of Kairserslautern, Technical University of Valencia. The main objective of the project is to develop the enabling technology and infrastructure required to effectively use the most advanced techniques developed for real-time applications with flexible scheduling requirements, in embedded systems design methodologies and tools, providing the necessary elements to target reconfigurable processing modules and reconfigurable distributed architectures. The approach to achieve this main objective is to integrate advanced flexible scheduling techniques directly into an embedded systems design methodology, covering all the levels involved in the implementation, from the OS primitives, through the middleware, up to the application level. This will be achieved by creating a contract model that specifies which are the application requirements with respect to the flexible use of the processing resources in the system, and also what are the resources that must be guaranteed if the component is to be installed into the system, and how the system can distribute any spare capacity that it has, to achieve the highest usage of the available resources.

http://www.frescor.org/

ACTORS - Adaptivity and Control of Resources in Embedded Systems

ACTORS is a three-year research project within the European Commission's 7th Framework Programme involving the following ArtistDesign partners: Lund University, Scuola Superiore Sant'Anna, and Technical University of Kairserslautern. It addresses design of resource-constrained software-intensive embedded systems with high requirements on adaptivity and efficiency. Three techniques are combined: virtualization, feedback control, and data-flow programming models. Virtualization techniques such as reservation-based scheduling provide spatial and temporal separation of concerns and enforce dependability and predictability. Reservations can be composed, are easier to develop and test, and provide security support. Using feedback-based resource management, the resource allocation is based on a comparison of the actual resource utilization of, e.g., a set of activities or tasks, with the desired resource utilization. The difference is used for deciding how the resources should be dynamically allocated. Feedback control makes it possible to deal with uncertainties and variations in a controlled way and provides adaptivity to on-line changes in objectives, external conditions and use cases. By combining feedback control with resource reservations it is possible to handle incorrect reservations, reclaim and redistribute unused resources, and adjust to dynamic changes in resource requirements.

http://www.actors-project.eu/index.php/

PREDATOR - Design for predictability and efficiency

PREDATOR is a three-year focused-research project within the European Commission's 7th Framework Programme on Research, Technological Development and Demonstration. The following ArtistDesign partners are involved: Saarland University, Scuola Superiore Sant'Anna, Technische Universität Dortmund, and University of Bologna. The project is concerned with embedded systems that are characterised by efficiency requirements on the one hand and critical constraints on the other. This combination of requirements typically occurs in application domains such as automotive, aeronautics, multi-media and industrial automation.

http://www.predator-project.eu/

THREAD Spanish project

The following ArtistDesign partners are involved: Technical University of Madrid, University of Cantabria, and Technical University of Valencia.

http://polaris.dit.upm.es/%7Estr/proyectos/thread/

WASP - Wirelessly Accessible Sensor Populations

The WASP project (Coordinated by Philips Research) aims at narrowing the mismatch between research at the application level and the node and network level in a sensor network, by covering the whole range from basic hardware, sensors, processor, communication, over the packaging of the nodes. The emphasis in the project lays in the self-organisation and the services, which link the application to the sensor network.

http://www.wasp-project.org/

iLAND – m**I**dd**L**ew**A**re for deterministic dynamically reconfigurable **N**etworke**D** embedded systems

iLAND is a European project of Call 1 of ARTEMIS JTU. The technical coordination is carried by University Carlos III de Madrid. The high level objective of iLAND is to develop enabling technology and infrastructure of a modular component-based middleware for in-building Infrastructured NES that have strong needs for deterministic dynamic functional composition and reconfiguration. The goal is to improve system flexibility, scalability, and composability. Also, maintainability will be improved since spontaneous reconfiguration of the system will be supported. This will enable dynamic functionality reconfiguration (i.e., new node activation, removal of crashed or damaged nodes and reallocation of functionality, etc.).

http://www.iland-artemis.org

-- Changes wrt Y1 deliverable --

No changes with respect to Year 1.

4. Internal Reviewers for this Deliverable

- **Gianluca Franchino** (University of Pavia)
- **Luis Almeida** (University of Porto, Portugal)