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

Activity - Progress Report for Year 1

Scheduling and Resource Management

Clusters:

Operating Systems and Networks

Activity Leader:

Professor Alan Burns (University of York, UK)

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

Policy Objective (abstract)

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

Versions

number	comment	date
1.0	First version delivered to the reviewers	December 19th 2008

Table of Contents

1. Overview of the Activity

1.1 ArtistDesign Participants and Roles

Professor Alan Burns (University of York - UK)
The real-time systems research group at the University of York contributes research on advanced scheduling and resource management policies.

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

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

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

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

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

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

Dr Stylianos Mamagkakis (IMEC)
The 'Runtime resource management for MPSoC' team at the Nomadic Embedded System division in IMEC is focusing on task scheduling, data storage and access methodologies to improve the performance and energy consumption of dynamic software applications, running on MPSoC platforms.

1.2 Affiliated Participants and Roles

Professor Alfons Crespo (Technical University of Valencia – Spain, Affiliated to Cantabria)
The team at the Technical University of Valencia is involved in providing real-time memory management OS support, and real-time kernel virtualization.

Professor Marisol García Valls (Carlos III University of Madrid - Spain, Affiliated to Cantabria)
The team at the Carlos III University of Madrid works on memory-based QoS management techniques to provide support for predictability in Real-Time Java middleware.

Professor Alejandro Alonso (Technical University of Madrid – Spain, Affiliated to Cantabria)
The team at the Technical University of Madrid investigates on integrated resource management policies with emphasis on adaptability.

Professor Lucia Lo Bello (Technical University of Catania - Affiliated to Pisa)
The team at the University of Catania works on QoS-oriented scheduling and management of communication and processing elements in embedded platforms, including energy-aware solutions.

Professor Pau Martí (Technical University of Catalonia – Affiliated to Lund)
The team at the Technical University of Catalonia work on the integration of feedback control and resource management techniques to provide adaptability to changing conditions on both resource and applications demands.

Professor Tullio Facchinetti (University of Pavia - Affiliated to the Scuola Superiore Sant'Anna)
The University of Pavia consider new methodologies for integrating overload management techniques with energy-aware strategies, in the context of small embedded systems for battery operated devices.

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

Professor Hermann Härtig (University of Dresden – Affiliated to Kaiserslautern)
The team at the University of Dresden are involved in building micro-kernel- and hypervisor-based systems as experimentation platforms.

1.3 Starting Date, and Expected Ending Date

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

1.4 Policy Objective

The main objective of this activity is the provision of models of embedded platform resources and policies, and the necessary analysis for undertaking the run-time scheduling of these resources and policies. A key scientific challenge is to link this resource-centred analysis with models of the application (and their resource usage policies) and the performance profiles of

the hardware platform itself. Issues of temporality, safety, reliability and security can only be effectively addressed by an integration of these various abstract views of the overall system.

Seven promising approaches for providing this integration are:

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

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

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

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

1.5 Background

The platforms on which the next generation of embedded systems will be implemented will be radically different from those used in the current generation. The scale, performance, scope and applicability are all subject to significant enhancement. This presents the application developer and systems engineer with a number of fundamental challenges. At the centre of these challenges is the (effective) management of the platform's resources. Such platforms are likely to be multi-core (64 soon and 200+ by 2010); involve buses and networks of various capabilities and speeds (both off-chip and on-chip, i.e. NoCs); memories of various speeds; include specialised components such as MEMS, ASICs, DSPs, and ASIPs; are linked to a wide variety of sensors and actuators; are embedded in systems powered by batteries (for mobile applications); include areas of FPGA (which are capable of dynamic reprogramming); and may have input/output links to global web-based information systems (for cyber-physical systems). Applications will be multi-resource and configurable. They will want to make dynamic modifications to their behaviour to support adaptability and environmental change. For example, the level of parallelism may alter at run-time and lead to re-evaluation of how this parallelism is delivered, e.g. by a subset of the cores, by application specific processing elements or by reprogramming an area of FPGA.

The main objective of this activity is to investigate how this wide variety of platform resources can be abstracted, modelled and managed, and application-specific resource allocation policies defined. At run-time, near optimal resource usage is desirable, but so are levels of protection for high integrity applications and those that have security constraints. Effective run-

time scheduling of multi-resource platforms is not currently achievable; new methods will need to be developed.

1.6 Technical Description: Joint Research

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

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

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

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

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

Dynamic memory management has been systematically avoided in real-time systems. One of the main reasons for this is the absence of deterministic allocators. Recently a new algorithm for dynamic memory allocation (TLSF) that solves this problem of the worst case bound whilst maintaining the efficiency of the allocation and deallocation operations has become available. This allows the reasonable use of dynamic memory management in real-time applications and permits consideration of dynamic memory as a first-class resource which can be used jointly with other resources in the schedulability of embedded systems. This integration of memory management and other resources is likely to develop over the next 18 months.

1.7 Problem Tackled in Year 1

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

- It will cover all forms of processing devices from single processors to multicores, FPGAs etc.
- It will cover all forms of communications including NoCs – although networking itself is covered by another activity within this cluster.

- It will cover all specialised components; ASIC etc – although platform issues are covered more in another cluster.
- It will cover system-wide resources such as energy.
- It will cover all forms of (offline) verification including simulations, analysis (scheduling analysis), and various forms of model checking – although this later topic is covered more in another cluster.
- It will cover online resource management (including control issues) to achieve openness, adaptability and fault tolerance.

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

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

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

2. Summary of Activity Progress

2.1 Technical Achievements

Towards a taxonomy of resource usage (All partners, lead by York)

As indicated in the previous section the key problem addressed during the first 18 months of this activity is to produce a taxonomy of resource usage. To this end meetings have taken place (including a cluster-wide meeting in Pisa) at which the framework for the taxonomy has been discussed and agreed. Work on the taxonomy will continue through the second year and a web-based report will be produced and be made available on the Artist public web site. There are of course many ways to structure a review such as this one into scheduling and resource management. Here the resources themselves are used to provide an initial decomposition. The following resource classes are deemed to be in scope: Processing resources, Communication resources, Specialised hardware, Memory/storage components, Data and software services and Power/energy and space. Each class has been subdivided into a number of resource types, for example the Specialised hardware class contains: MEMS, ASICs, DSPs, ASIPs, GPUs, Coders/decodes of various types, Sensors and Actuators.

Flexible scheduling framework (Cantabria, York, Pisa, Kaiserslautern, Valencia, Aveiro, and several academic and industrial partners)

An architectural model of a flexible scheduling framework had been developed in previous years as part of the FIRST and FRESCOR EU projects and has been extended in the current reporting period. The framework is capable of handling multiple concurrent activities with different criticality and timing in the same system, integrating the management of different kinds of resources such as processors, networks, memory, energy, and shared objects with time protection. The framework is independent of the underlying implementation, and can run on different underlying scheduling strategies. It is based on establishing service contracts that represent the complex and flexible requirements of the applications and which are managed by the underlying system to provide the required level of service.

Based on the experience gained during previous years the contract model has been updated and its API has been redesigned. The main changes made to the contract model have been:

- Updated energy management API to add operations to get more information from the system.
- New disk-bandwidth management services. A very simple module has been added to the framework to explore the possibility of handling contracts related to disk bandwidth.

Work has been performed in cooperation with Thales Communications France (TCF) to integrate the contract-based scheduling framework with a component-based framework. The component based technology is the microCCM framework that implements the component-container model with an infrastructure that is independent of CORBA.

Multi-resource Scheduling on Multicore Platforms (NXP, TUKL)

NXP and TUKL continued work started during ARTIST2 on integrating real-time scheduling and cache management on multiprocessor platforms. They carried out experiments to study cache behaviour on the actual platform and formulated a number of scenarios with increasing complexity. A joint PhD student is supported by this activity.

Memory arbitration on Heterogeneous Multicore platforms (NXP)

NXP also performed work regarding memory arbitration in a heterogeneous multiprocessor environment, applying techniques that have been developed for processor sharing by other research partners. This activity has lead to implementations in NXP products.

Flexible control on low cost microcontrollers (University of Aveiro and Technical University of Catalonia)

In recent years, approaches to control performance and resource optimization for embedded control systems have been receiving increased attention. Most of them focus on theory, whereas practical aspects have been omitted. Theoretical advances demand flexible real-time kernel support for multitasking and preemption, thus requiring more sophisticated and expensive software/hardware solutions. On the other hand, embedded control systems often have cost constraints related with mass production and strong industrial competition, thus demanding low-cost solutions. This work shows that these conflicting demands can be softened and that a compromise solution can be reached. It is advocated that recent research results on optimal resource management for control tasks can be implemented on simple multitasking preemptive real-time kernels targeting low-cost microprocessors, which can be easily built in-house and tailored to actual application needs. The experimental evaluation shows that significant control performance improvement can be achieved without increasing hardware costs. A joint publication to IEEE Transactions on Industrial Electronics has resulted from this work.

Schedulability analysis for CAN-based control applications with dynamic bandwidth management (Technical University of Catalonia)

This work focuses on the schedulability analysis for control messages when networked control loops built on top of the Controller Area Network (CAN) are dynamically allocated bandwidth in terms of their controlled plants' dynamics. The bandwidth allocation policy is theoretically described by an optimization problem and practically solved by the distributed bitwise arbitration of CAN messages when message identifiers, i.e., priorities, reflect control applications demands. This poses the problem of assessing whether the set of real-time messages will meet their deadlines regardless of run-time priority changes. This is solved by schedulability analysis based on recent results on worst-case response time techniques for real-time CAN applications. The analysis ends up with the schedulability test for this type of applications.

Profiling and Analysis (IMEC, Univ. Complutense Madrid, and Democritus Uni. of Thrace)

A profiling and analysis methodology has been produced to collect software metadata information. The metadata characterize the dynamic data access and storage patterns which are allocated in the heap memory and have a dynamic size and lifetime. This software metadata information is used to steer dynamic data type (DDT), dynamic memory management (DMM) and direct memory access (DMA) optimization methodologies at the system level. The optimizations achieved performance gains of 43%.

Sensitivity Analysis (Pisa)

In the first year, the Scuola Superiore Sant'Anna of Pisa has addressed the issue of providing sensitivity analysis for fixed-priority real-time systems, with the objective of deriving a feasibility

region that can be used to improve the design of real-time embedded systems. The schedulability region has been used to take resource constraints into account in finding the task deadlines that optimize the control performance.

Optimal period selection and scheduling for embedded controllers (ULUND, Pisa, Linköping Univ))

When several digital controllers should execute on the same platform, a key question is how to distribute the computing resources among the control loops. For controllers based on periodic sampling, this boils down to selecting sampling periods for the set of controllers. In this work, ULUND and SSSA have considered a static design problem, where the objective is to assign sampling periods in order to minimize the aggregate cost (performance index) of the controllers. For controllers, the delay between sampling and actuation has a great impact on the performance. For this reason, it is important to include a model of the delay in the design problem. An approximate response-time analysis has been developed that allows estimating the delay for a controller under fixed-priority scheduling. Assuming cost functions that are linear in period and delay, the approximate analysis allows giving analytical expressions for the optimal sampling periods. A case study with random second- and third-order plants is used to evaluate the improvements compared to methods that do not take the delay into account.

A similar problem has also been studied by ULUND and Linköping University for distributed systems. The system consists of a number of nodes, connected by a communication bus. The nodes and the bus may be scheduled statically or dynamically (using priorities). Again, the objective is to minimize the combined cost of the controllers in the application. The design problem is solved using a genetic algorithm that decides the execution pattern of the control tasks. In the dynamic scheduling case, simulation is used to estimate the average delay and jitter of the control tasks. The Jitterbug toolbox from Lund is used to evaluate the control performance, taking the delay and jitter into account. A case study with several benchmark plants is used to evaluate the performance compared to previous, heuristic design methods.

Sporadic Event-Based Control (ULUND)

Normally, controllers are designed assuming equidistant (periodic) sampling. This simplifies the design process greatly, since the sampled plant description becomes a linear time-invariant (LTI) discrete-time system. From a computing or network point of view, it makes sense to only sample or control when something significant has occurred in the system. In this line of work, ULUND has investigated several aspects of event-based control, including analysis of how delay, jitter and measurement noise affect the performance of the control loops, scheduling of multiple sporadic controllers on a shared communication medium, considering TDMA, FDMA and CSMA medium access methods, analysis of limit cycles in event-based control, implementation work on event-based PI control, and event-based state estimation. Part of this work involves collaboration with UPC and KTH.

Real-Time Scheduling on Multicores (Polytechnic Institute of Porto, Portugal)

Real-time scheduling on a single processor has enjoyed great successes during the recent four decades, starting with the development of rate-monotonic which was used in the computer that put the first man on the moon. Those decades of developments of scheduling theory for a single processor have made it possible to (i) schedule tasks with few preemptions and small time-complexity of dispatching, (ii) schedule implicit-deadline tasks to meet deadlines even at high processor utilization (69% and 100%), (iii) schedule arbitrary-deadline tasks and (iv) schedule tasks with the aforementioned properties but also to share data between tasks. Unfortunately such results are not available for the computer platform that matters the most

today: the multiprocessor implemented on a single chip, also called multicore. Therefore, researchers at Polytechnic Institute of Porto work on the development of those results as part of a national project called RESCORE.

Dynamic run-time adaptability (Polytechnic Institute of Porto, Portugal)

The growing complexity and dynamism of many embedded application domains (including consumer electronics, robotics, automotive and telecommunications), makes it increasingly difficult to react to load variations and adapt the system's performance in a controlled fashion within an useful and bounded time. This is particularly noticeable when intending to benefit from the full potential of an open distributed cooperating environment, where service characteristics are not known beforehand and tasks may exhibit unrestricted QoS inter-dependencies.

In this context, in this year, the Polytecnhic Institute of Porto has addressed the increased complexity of negotiating service provisioning for task sets which exhibit unrestricted QoS dependency relations. The proposed approach is based in anytime algorithms, which can be interrupted at any time and still provide a solution and a measure of its quality, expected to improve as the run time of the algorithms increases. This flexibility in the algorithms' execution times enables the system to timely adapt to the changing environmental conditions of dynamic open real-time systems.

This research activity also addressed the scheduling of tasks that share resources and exhibit precedence constraints, without a previous complete knowledge about their behaviour. This was achieved by integrating the concept of bandwidth inheritance with the greedy capacity sharing and stealing policy of CSS. Rather than trying to account borrowed capacities and exchanging them later in the exact same amount, the new approach greedily exchanges extra capacities as early, and not necessarily as fairly, as possible. The achieved results suggest that the approach effectively minimises the impact of bandwidth inheritance on blocked tasks.

Integrated Memory and Communications Management based on RTSJ (Universidad Carlos III de Madrid, Universidade Aveiro)

RTSJ-based extensions to non garbage collected memory have been made by UC3M and applied to the remote invocation model of Java. The communication model has been ported on top of Flexible Time-Triggered Ethernet of the University of Aveiro. This effort goes in the direction of achieving a predictable integrated resource management by combining memory and network management.

Assessment of the IEEE 802.15.4 GTS scheduling and allocation mechanism for real-time Wireless Sensor Networks (University of Catania)

The IEEE 802.15.4 protocol is generally considered one of the most suitable protocols for wireless real-time communication. This is for two main reasons. First, it features a more robust coding technique compared with other wireless standards, such as the IEEE 802.11 family. Second, it supports both contention access transmission and contention-free transmission using Guaranteed Time Slots (GTS). Previous research already proved that, in presence of no transmitting errors, a bounded delay can be achieved when the star and cluster-tree toplogies are used for contention-free GTS transmission. The paper [CT1] investigates the limitations of the cluster tree topology for the implementation of large-scale real-time WSNs. Simulation results obtained using the Omnet++ tool showed that two problems, namely low scalability and unbalanced energy consumption, limit the suitability of the IEEE 802.15.4 for such kinds of

networks. The paper analyzes the causes of these problems and proposes viable alternatives to address them. Simulation results in terms of throughput and delay are presented.

Combined Energy and QoS management in WSNs through topology control (University of Catania)

Energy efficiency and QoS provisioning are two main requirements for Wireless Sensor Networks (WSNs). Unfortunately, such requirements clash with each other, thus most of the WSN protocols are able to effectively achieve only one of them. Topology control protocols are used in wireless ad-hoc networks to decrease the duty cycle of nodes while maintaining the network connectivity. Such protocols can exploit the redundancy of densely deployed WSNs to prolong the network lifetime. They can be used in conjunction with QoS-enabled routing protocols to achieve QoS support and energy efficiency at the same time. However, classical approaches, e.g., GAF or SPAN, are not suitable for QoS enabled or real-time systems, as it is not possible to exactly quantify the delay they introduce and the reduction in routing fidelity they cause. In [CT2] is presented a topology control protocol that, in addition to highly increase the network lifetime, has two important properties, namely bounded introduced delay and guaranteed routing fidelity, that make it suitable for QoS-enabled WSNs.

These features are obtained thanks to the combination of time division channel access mechanism with a cellular radio architecture. This paper describes in detail the protocol and analyzes its behaviour in terms of energy efficiency as a function of the topology and configuration parameters. This enable the estimation of the trade-off between power consumption and data delivery speed at design time. Finally, it presents experimental results obtained through the ns-2 simulator, which confirm the analytical results on energy-consumption and assess the effect of the proposed topology control mechanism on the routing performance. In particular, a comparison between the performance of a well-know routing protocol, i.e., the SPEED protocol, when it is used with or without the proposed topology management mechanism, respectively, show that the proposed topology control protocol is able also to increase the overall network capacity and so the real-time performances.

2.2 Individual Publications Resulting from these Achievements

University of York

R.I. Davis, A. Burns (2008), An Investigation into Server Parameter Selection for Hierarchical Fixed Priority Pre-emptive Systems. Proceedings of Real-Time and Network Systems, RTNS.

R.I. Davis, A. Zabos, A. Burns (2008), Efficient Exact Schedulability Tests for Fixed Priority Real-Time Systems. IEEE Transactions on Computers, Vol. 57, No. 9, pp. 1261-1276.

A. Zuhily, A. Burns (2008), Exact Response Time Scheduling Analysis of Accumulatively Monotonic Multiframe Real Time Tasks. 5th International Colloquium on Theoretical Aspects of Computing (ICTAC).

A. Zuhily, A. Burns (2008), Exact Scheduling Analysis of Accumulatively Monotonic Multiframe Tasks Subjected to Release Jitter and Arbitrary Deadlines. 13th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA).

A. Zuhily, A. Burns (2008), Exact Scheduling Analysis of Non-Accummulatively Monotonic Multiframe Tasks. Proceedings of Real-Time and Network Systems, RTNS.

Y. Chu, A. Burns (2008), Flexible hard real-time scheduling for deliberative AI systems. Real-Time Systems Journal

S. Baruah, A. Burns (2008), Quantifying the sub-optimality of uniprocessor fixed-priority scheduling. Proceedings of Real-Time and Network Systems, RTNS.

A. Burns, S. Baruah (2008), Sustainability in Real-time Scheduling, Journal of Computing Science and Engineering. pp7—97, Vol. 2 , No 1.

TUKL

Raphael Guerra, Gerhard Fohler, "A Gravitational Task Model for Target Sensitive Real-Time Applications", ECRTS08 - 20th Euromicro Conference on Real-Time Systems, Prague, Czech Republic. 2008.

Universidad de Cantabria

Patricia López Martínez, Julio Medina, & José María Drake. "Real-Time extensions to "Deployment and Configuration of Component-based Distributed Applications". OMG Workshop on Distributed Object Computing for Real-Time Embedded Systems, Washington, DC, USA, July, 2008.

Patricia López Martínez, Julio L. Medina, Pablo Pacheco and José M. Drake. "Ada-CCM: Component-based Technology for Distributed Real-Time Systems" 11th International Symposium on Component Based Software Engineering (CBSE-2008) Karlsruhe, Germany October 2008.

J.L. Gilbert, O. Hachet, J. Chauvin, P. López, J.M. Drake, M. González Harbour, "Integration of Flexible Real-Time Scheduling Services in a Lightweight CCM-Based Framework". OMG Workshop on Distributed Object Computing for Real- Time Embedded Systems, Washington, DC, USA, July, 2008.

Pisa

Enrico Bini, Marco Di Natale, and Giorgio Buttazzo, "Sensitivity Analysis for Fixed-Priority Real-Time Systems", Real-Time Systems, Vol. 39, No. 1-3, pp. 5-30, August 2008.

Yifan Wu, Enrico Bini, and Giorgio Buttazzo, "A Framework for Designing Embedded Real-Time Controllers", Proceedings of the 14th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA 2008), Kaohsiung, Taiwan, August 25-27, 2008.

Lund University

Anton Cervin, Erik Johannesson: "Sporadic Control of Scalar Systems with Delay, Jitter and Measurement Noise". In Proc.17th IFAC World Congress, Seoul, Korea, July 2008.

Anton Cervin, Toivo Henningsson: "Scheduling of Event-Triggered Controllers on a Shared Network." In Artist2 invited session at Proc. 47th IEEE Conference on Decision and Control, Cancun, Mexico, December 2008.

Toivo Henningsson: "Recursive State Estimation for Linear Systems with Mixed Stochastic and Set-Bounded Disturbances." In Proc. 47th IEEE Conference on Decision and Control, Cancun, Mexico, December 2008. To appear.

Polytechnic Institute of Porto

B. Andersson, K. Bletsas, Sporadic Multiprocessor Scheduling with Few Preemptions. Euromicro Conference on Real-Time Systems, Prague, Czech Republic, July, 2008.

B. Andersson, K. Bletsas, S. Baruah, Scheduling Arbitrary-Deadline Sporadic Tasks on Multiprocessors, 29th IEEE Real-Time Systems Symposium, Barcelona, Spain, December , 2008.

B. Andersson, The Utilization Bound of Uniprocessor Preemptive Slack-Monotonic Scheduling is 50%, 23rd Annual ACM Symposium on Applied Computing, Vila Galé in Fortaleza, Ceará, Brazil, 2008.

B. Andersson, Schedulability Analysis of Generalized Multiframe Traffic on Multihop-Networks Comprising Software-Implemented Ethernet-Switches, 16th International Workshop on Parallel and Distributed Real-Time Systems (WPDRTS '08), April 14, 2008, Miami, Florida, USA.

B. Andersson, Static-Priority Preemptive Multiprocessor Scheduling with Utilization Bound 38%, 12th International Conference on Principles of Distributed Systems, Luxor, Egypt, December 15-18, 2008.

B. Andersson, Uniprocessor EDF Scheduling with Mode Change, 12th International Conference on Principles of Distributed Systems, Luxor, Egypt, December 15-18, 2008.

Luís Nogueira, Luís Miguel Pinho. "Dynamic QoS Adaptation of Inter-Dependent Task Sets in Cooperative Embedded Systems". In Proceedings of the 2nd ACM International Conference on Autonomic Computing and Communication Systems, ACM Press, 2008.

Luís Nogueira, Luís Miguel Pinho. "Handling QoS Dependencies in Distributed Cooperative Real-Time Systems". In IFIP Distributed Embedded Systems: Design, Middleware and Resources, Springer, 2008.

Luís Nogueira, Luís Miguel Pinho. "Shared Resources and Precedence Constraints with Capacity Sharing and Stealing". In Proceedings of the 22th IEEE International Parallel and Distributed Processing Symposium (IPDPS 2008), IEEE Computer Society Press, 2008.

Luís Nogueira, Luís Miguel Pinho. "Capacity Sharing and Stealing in Dynamic Server-based Real-Time Systems". In Proceedings of the 21th IEEE International Parallel and Distributed Processing Symposium (IPDPS 2007), IEEE Computer Society Press, 2007.

University of Catalonia

M. Velasco, P. Martí, J.M. Fuertes (2008), Schedulability analysis for CAN-based control applications with dynamic bandwidth management, Technical University of Catalonia, Automatic Control Department, Research Report ESAII-RR-08-04, November 2008.

NXP Semiconductors

L.Steffens, M. Agarwal, P. van der Wolf, “Real-Time Analysis for Memory Access in Media Processing SoCs: A Practical Approach”, ECRTS08 - 20th Euromicro Conference on Real-Time Systems, Prague, Czech Republic. 2008.

University Carlos III de Madrid

P. Basanta-Val, M. García-Valls, and I. Estévez-Ayres. Simplifying the Dualized Threading Model of RTSJ Proc. of the 11th IEEE International Symposium on Object/component/service-

oriented Real-time distributed Computing, ISORC 2008. May 5 - May 7, 2008. Orlando, Florida, USA.

P. Basanta-Val, M. García-Valls, I. Estévez-Ayres y J. Fernández-González. Integración de capacidades de multiplexación en el conjunto de subprotocolos JRMP. IEEE América Latina, ISSN 1548-0992. In Spanish.

University of Catania

E. Toscano, L. Lo Bello, “On the Use of IEEE 802.15.4 for Real-Time Wireless Sensor Networks”, In Proceedings of the 7th International Workshop on Real-Time Networks, RTN’08, Prague, Cech Republic, July 2008.

E. Toscano, L. Lo Bello, “A Topology Management Protocol with Bounded Delay for Wireless Sensor Networks”, In Proceedings of the 13th IEEE International Conference on Emerging Technologies and Factory Automation, ETFA’08, Hamburg, Germany, Sep. 2008, pp. 942 - 951, ISBN 978-1-4244-1505-2.

2.3 Interaction and Building Excellence between Partners

The partners involved in this activity are part of an international community that meets regularly at the main conferences for the disciplines covering scheduling and resource management. These conferences include the IEEE International symposium RTSS and the European conference on real-time issues, ECRTS. Interactions also take place as part of other funded EU projects such as FRESCOR. Finally ArtistDesign specific meeting have taken place, for example a meeting in Pisa in October 2008. At this meeting a framework for the taxonomy on resource usage was debated and agreed upon (in outline).

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

ULUND, SSSA, TUKL, Ericsson, and Evidence collaborate on feedback-based resource scheduling for multimedia terminals within the EU STREP project ACTORS (Adaptivity and Control of Resources in Embedded Systems).

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

UC3M and U. Aveiro collaborate actively in the integration of the work on resource-management based on RTSJ memory model and service-based application on top of real-time communications support (FTT).

2.4 Joint Publications Resulting from these Achievements

S.Perathoner, E. Wandeler, L.Thiele, A.Hamann, S.Schliecker, R.Henia, R.Racu, R.Ernst, M. G. Harbour, "Influence of different abstractions on the performance analysis of distributed hard real-time systems". Journal on Design Automation for Embedded Systems, Springer, April, 2008, ISSN: 0929-5585.

Ricardo Marau, Pedro Leite, Manel Velasco, Pau Martí, Luis Almeida, Paulo Pedreiras and Josep M. Fuertes (2008), Performing Flexible Control on Low-Cost Microcontrollers Using a Minimal Real-Time Kernel, In IEEE Transactions on Industrial Informatics, vol. 4, n.2, May 2008.

Alexandros Bartzas, Miguel Peon-Quiros, Stylianos Mamagkakis, Francky Catthoor, Dimitrios Soudris, Jose Manuel Mendias: Enabling run-time memory data transfer optimizations at the system level with automated extraction of embedded software metadata information. ASP-DAC, 434-439. 2008.

Iria Estévez-Ayres, Marisol García-Valls, Luís Almeida, and Pablo Basanta-Val. *Solutions for Supporting Composition of Service-Based Real-Time Applications* Proc. of the 11th IEEE International Symposium on Object/component/service-oriented Real-time distributed Computing, ISORC 2008. May 5 - May 7, 2008. Orlando, Florida, USA.

2.5 Keynotes, Workshops, Tutorials

Keynote:
Karl-Erik Årzén (ULUND): "Adaptivity in Embedded Systems – Why, What and How" *Workshop on Adaptive and Reconfigurable Embedded Systems, St Louis, April 21, 2008.*

Workshop:
Multicores: Theory and Practice, *Kaiserslautern, October 27. The workshop was organized by ArtistDesign and ACTORS in order to discuss schedulability issues related to multicore platforms. Six invited presentations were given.*
http://www.artist-embedded.org/artist/Overview,1501.htm

Workshop:
Operating Systems and Networks, *Pisa, Italy – October 2-3, 2008.* The purpose of the workshop was to refine the research objectives of the cluster on Operating Systems and Networks and coordinate the collaboration among the different groups. Discussed topics included: architecture effects on worst-case execution times, taxonomy of resources, real-time networks, and real-time and control issues.

Organizers:

- Giorgio Buttazzo - Scuola Superiore Sant'Anna, Italy
- Alan Burns – University of York, UK
- Luis Almeida, Univ. of Aveiro, Portugal

3. Milestones, and Future Evolution

3.1 Problem to be tackled over the next 12 months (Jan 2009 – Dec 2009)

The main cluster-wide problem to be tackled during the next 12 months is to complete the taxonomy of resource usage that has started during the first year. This report will be published on the Artist public web site. The objective of the review is to capture the policies, analysis techniques, and modeling approaches that are appropriate for resource management for a wide class of embedded systems. Both offline verification and online control will be covered. A broad range of resources will be included (see summary of first year activities).

In general policies provide for effective resource usage; analysis provides for predicting system behaviour using simulation, scheduling analysis, measurement, testing or model checking. Models provide the link with the application and facilitate subsystem composition using, for example, time triggered and event-triggered work flows, and static and dynamic usage patterns.

The challenge for the research community is to move from single processor platforms to multiprocessor, multi-core, FPGA etc; and to integrate various resource and abstraction views of the overall system by:

- Integrate policies
- Integrate analysis
- Integrate models

In doing so, various component compositions must be accommodated including static and dynamic interactions, and peer-to-peer and hierarchical architectural structures. The taxonomy will provide links to the key material for this future work. Once produced the taxonomy will form a resource for framing and informing the future work on this Activity.

A key problem to be tackled in the next 12 months (and beyond) will be the international effort at deriving effective scheduling methods for multiprocessor (eg MPSoC). This remains a key challenge; with the lack of usable algorithms and verification techniques being a real inhibitor to the deployment of these new platforms. Both theoretical results and practical algorithms are needed. The former may assume pure homogeneous processors linked via a predicable communication medium; the latter must consider heterogeneous systems with multiple levels of cache and best-effort NoCs.

Other work will continue to address scheduling problems such as effective EDF scheduling and implementation, FPGA platforms, memory modelling, energy and power control, contract-based allocation of resources, and adaptive dynamic behaviour. These issues will also be addressed in the scope of traffic scheduling in communication systems, together with the networks activity, particularly concerning communication switches and complex network topologies, e.g., wireless sensor networks and multi-segment hybrid networks. Many of the key research teams in this domain of scheduling and resource usage are members of ArtistDesign. These teams are expected to continue to influence the directions of future research.

3.2 Current and Future Milestones

- Produce a taxonomy of resource usage – to be completed during year 2.
- Extend the use of hierarchical and contract-based scheduling to multi-resource systems.
- Produce effective scheduling and placement algorithms for multiprocessor systems.
- Extend sensitivity analysis to EDF and multiprocessor systems.
- Produce mode change algorithms suitable for multiprocessor systems.
- Determine an effective way of undertaking (static) architectural tradeoffs.
- Determine an effective way of undertaking (dynamic) adaptive resource management (making use of feedback techniques from the control environment).
- Define a framework that can accommodate multiple time-frames within a single system and facilitate hierarchical scheduling, cascade control and other means of separating temporal concerns.

3.3 Main Funding

In addition to the specific funds from the ArtistDesign NoE, the main sources of funding are:

- JEOPARD – EU funded project on parallel real-time Java, consortium includes Vienna and York.
- HIJA – High Integrity Java – EU funded project in which the following partners are involved: University of York, University of Madrid.
- Javamen – Java on FPGA platforms – UK national DTI project involving the University of York.
- eMuCo – EU funded project aimed at developing a platform for future mobile devices based on multi-core architecture, consortium included York, the Ruhr-Universität Bochum, Technische Universität Dresden and Politehnica University of Timisoara, as well as industrial partners, such as Infineon, Telelogic, ARM and GWT-TUD.
- Reflect - Reflection Mechanisms in Real-Time Embedded Systems – Portuguese funded FCT project involving the Polytechnic Institute of Porto.
- CooperatES - QoS-Aware Cooperative Embedded Systems – Portuguese funded FCT project involving the Polytechnic Institute of Porto.
- RESCUE - REliable and Safe Code execUtion for Embedded systems – Portuguese funded FCT project involving the Polytechnic Institute of Porto.
- The FRESCOR EU IST project, in which the following ArtistDesign partners are involved: University of Cantabria, University of York, Scuola Superiore Santa Anna, Technical University of Kairserslautern, Technical University of Valencia – www.frescor.org
- THREAD Spanish project, in which the following ArtistDesign partners are involved: Technical University of Madrid, University of Cantabria,Technical University of Valencia
- ASSERT - Automated proof based System and Software Engineering for Real-Time applications. EU funded project. Main objective is to improve the system-and-software

development process for critical embedded real-time systems in the Aerospace and Transportation domains.

- MADEJA: French funded (Region Rhône-Alpes) in which VERIMAG was involved for studying quantitative prediction of dynamic memory allocation for RTSJ scoped-memory mechanism.
- MOSART IST-215244 Project: Mapping Optimization for Scalable multi-core ARchiTecture. ArtistDesign partners involved: Interuniversitair Micro-Elektronica Centrum (IMEC) vzw. and Kungliga Tekniska Hagskolan (KTH) http://www.mosart-project.org/
- MNEMEE IST-216224 Project: Memory maNagEMEnt technology for adaptive and efficient design of Embedded systems. ArtistDesign partners involved: Interuniversitair Micro-Elektronica Centrum (IMEC) vzw. and Technische Universiteit Eindhoven (TU/e) http://www.mnemee.org/
- RESCORE. Portuguese funded project aiming to develop real-time scheduling algorithms for multicores.
- HARTES - HArd Real-Time Ethernet Switching - Portuguese funded FCT project involving the University of Aveiro.
- OMP IST STREP – Incorporating QoS and Resource Management APIs in Open Media Platform standard OpenMax.
- FlexWARE - Catania is involved in the IST 7FP STREP project “Flexible Wireless Automation in Real-Time Environments”,aiming at the implementation of a novel platform for the support of real-time communication over Wireless Local Area Networks based on the IEEE 802.11standard.

4. Internal Reviewers for this Deliverable

Michael González Harbour (Universidad de Cantabria)

Peter Puschner (Technische Universitaet Wien)

Andy Wellings (University of York)