ARTIST 2
Network of Excellence

IST-004527 ARTIST2:
Embedded Systems Design

Activity Progress Report for Year 2

JPRA-Cluster Integration

Adaptive Resource Management for Consumer Electronics

Clusters:

Main Adaptive Real-time ART

Activity Leader:

Gerhard Fohler (Technische Universität Kaiserslautern)
http://www.eit.uni-kl.de/fohler/

Resource management for consumer electronics faces the challenges of highly variable resource demand, e.g., from MPEG encoded video streams, highly variable resource availability, e.g., wireless networks, shared CPUs, and strong requirements for resource efficiency due to cost considerations in mass markets. Violations of timeliness requirements both on a single resource as well as in the end-to-end chains lead to unacceptable quality impairments.

This requires attention to a number of issues in diverse areas of research. The activity addresses these issues by leveraging on the critical mass provided by ARTIST2 and in the ART cluster, with expertise ranging from flexible scheduling and infrastructure to wireless networks and media processing. Expertise outside the team is made accessible via industrial contacts to provide for realistic application requirements.

Table of Contents

1. Overview of the Activity

1.1 ARTIST2 Participants: Expertise and Roles

Activity Leader: Gerhard Fohler - Technische Universität Kaiserslautern
resource management framework, video streaming, wireless networks

Team Leader: Giorgio Buttazzo - Scuola Superiore S. Anna
scheduling mechanisms for temporal isolation

Team Leader: Michael Gonzalez Harbour - University of Cantabria
end-to-end timing analysis of distributed systems

Team Leader: Luis Almeida - University of Aveiro
bandwidth adaptation in networks

Team Leader: Eduardo Tovar - Polytechnic Institute of Porto
communication protocols for wireless networks

Team Leader: Alejandro Alonso - UP Madrid
adaptive techniques for QoS

1.2 Affiliated Participants: Expertise and Roles

Team Leader: Marisol García-Valls - U. Carlos III, Madrid
QoS management

Team Leader: Pau Martí - Universitat Politècnica de Catalunya
real-time and control

Team Leader: Paolo Gai - Evidence, SME
tools for resource management

Team Leader: Liesbeth Steffens - Philips Research
consumer electronics, video streaming, multi resource management

1.3 Starting Date, and Expected Ending Date

It started in October 2004, and will be complete when research topics of industrial relevance and example solutions have been identified and developed.

Expected ending date: March 2008

1.4 Baseline

In some application domains, such as multimedia, applications are very expensive in terms of resource consumption. In other applications domains, such as automotive, mobile telephony or even building automation, the resources are scarce and there is a growing pressure to integrate resources even further and optimize their use. In both cases, timeliness directly relates to user perceived quality, e.g., smoothness of the video stream. Furthermore, efficient resource usage is key issue not only for cost considerations, but also for competition on a feature bases: better resource usage – more features.

Both resource demands, e.g., MPEG-2 video streams, and resource availability, e.g., available bandwidth on wireless links, fluctuate rapidly and unpredictably; worst case assumptions will

lead to extreme over provisioning. Consequently, methods for adaptive resource management are required.

Trading resource usage (processing, communication and memory/storage, inter-device and intra-device) against offered output is known as QoS (Quality of Service). The different resources cannot be considered separately, interferences and inter-resource tradeoffs have to be taken into account because they affect the application output. The tradeoffs have to be made at different time scales, in order to match the time scales of the system dynamics.

1.5 Problem Tackled in Year2

The objectives of the activity in the second year crystallised around the following issues:

- **Application requirements.** Applications in consumer electronics, in particular video processing, show similar demands to classic real-time applications, e.g., the importance of meeting deadlines for the display for frames. The characteristics and requirements, however, diverge significantly: the amount of resources needed varies largely, e.g., due to frame sizes and types, and are diverse for resources. A small frame will require little bandwidth for transmission, but may induce high decoding times. Looking at the end-to-end delivery chain, diverse devices and their resource will impact the temporal behaviour of a stream. Semantic stream dependencies, as with MPEG encoding have a similar impact. In particular, we looked into the problem of the temporal impact of devices and demands of video streaming with input from the affiliated partners in consumer electronics. A workshop on requirements for adaptive resource management was held with the FRESCOR project.
- **Adaptive Management of Multiple Resources.** The resources typically used in end-to-end delivery of streams often exhibit fluctuating availability. Wireless networks, for example, are influenced by interference, mobility, or physical structures, which cannot be known before system deployment. A key challenge here is the provision of methods indicating bandwidth availability during system runtime. Ideally, predictions about availability in future time interval can be given with meaningful confidence, but notion of achievable Quality-of-Service will be useful already.

  Even on single devices, a number of resources will be interdependent. Thus the issue of multi resource management is important. In particular, modern multiprocessor architectures for media processing demand the integration of CPU scheduling and cache/memory handling for predictable timing.

- **QoS Middleware.** In addition to these effects on single devices, he fluctuating resource availability has to be dealt with on a system wide level as well, allowing the system to make QoS adaptation decisions in response to fluctuations in the nodes' service load, under the control of the user. Since complex dynamic real-time scenarios may prevent the possibility of computing optimal service configurations before execution, iterative refinement approaches with the ability to trade off deliberation time for the quality of the solution will be more practical. Another issue is the adaptation of distributed QoS negotiation policies the integration of this middleware with communication protocols.

1.6 Comments From Previous Review

1.6.1 Reviewers' Comments

The goal of this task was to define techniques for adaptive resource management according to QoS by identifying requirements from consumer electronics and automation.

The activity progress from months 1-6 and months 7-12 is now clear, but the funding is not defined. The partner relationship and collaboration and steps taken to build a network of excellence are now appearing. The industrial needs and experience are now showing perspectives.

The deliverable is accepted as it is.

1.6.2 How These Have Been Addressed

The funding of activities, apart from partners own resources is outlined in section 3.4.

2. Summary of Activity Progress

2.1 Previous Work

In the first period, the technical results were achieved in the following areas: video stream demand analysis, identification of scheduling algorithms and kernel mechanisms for stream adaptations based on integrated, flexible scheduling; adaptive resource management for network bandwidth management, multi resource management, in particular with respect to cache aware scheduling; middleware support for QoS management.

Furthermore, the ART cluster has been in active contacts with relevant industry to gather understanding of realistic requirements and to identify research topics and baselines relevant for industrial and academic research. Partners has been giving presentations at the Philips Software Conference – Real-time Workshop and had meetings with Nokia, Ericsson mobile platforms and Visual tools from Spain. The goal has been to go as far as possible towards the actual engineers for better understanding and prepare for a specific industry – academia workshop with selected participants.

2.2 Current Results

2.2.1 Technical Achievements / Outcomes / Difficulties encountered

Temporal Constraints for Video streaming

Philips and TUKL have studied temporal constraints of video streaming. As sources for the constraints we looked into semantic stream dependencies from MPEG decoding, as well as the temporal impact of devices and their resources in the end-to-end delivery chain of a stream. The work was carried out with industrial partners in the area. The results will feed into other activities in the cluster, in particular w.r.t. to scheduling and networking. http://rts.eit.uni-kl.de/research/mediaprocessing

Difficulty: The workshop for requirements analysis with industrial engineers could not be organized due to a major restructuring within the industrial partner.

Integrated real-time scheduling and cache management

Philips and TUKL have continued work on integrating real-time scheduling and cache management on multiprocessor platforms. To this end, we have carried out experiments to study cache behaviour on the actual platform and formulated a number of scenarios with increasing complexity. The work is being carried out by a joint PhD student.

Server Based Flexible Scheduling

Schedulability analysis techniques for server-based systems that can be used to schedule different kinds of flexible timing requirements, such as those needed to integrate control systems with multimedia activities. In particular, this work has been focused on hierarchical scheduling analysis and design techniques A further issue has been the dimensioning the parameters of a server for minimizing the average response time of the served activities. A statistical approach has also been addressed in order to compute the probability of missing a given deadline.. Partners were SSSA, Cantabria, TUKL of the cluster and the partners of the FIRST and FRESCOR EU STREP consortia. www.frescor.org http://rts.eit.uni-kl.de/research/adaptive-rts
http://cordis.europa.eu/fetch?CALLER=PROJ_IST&ACTION=D&RCN=62751&DOC=1&CAT=PROJ&QUERY=1158229719107

Adaptive resource management for networks

Work concerned the analysis of the achievable QoS guarantees in wireless networks. After defining a proper model for the main resources (i.e., CPUs, disk and network), a number of existing scheduling algorithms for the three types of resources have been analysed under fluctuating workload to evaluate their behaviour in terms of service guarantee. Then, the achievable end-to-end QoS guarantees have been investigated as a function of the guarantees provided by the underlying resource schedulers. Further activities dealt with network protocols to efficiently support dynamic bandwidth management with strict QoS guarantees in Ethernet-based networks, which is still an important networking technology in the field of distributed multimedia systems. A wireless time-token communication protocol that allows providing real-time guarantees for real-time messages and tune the allocated bandwidth according to the required QoS was developed. Aveiro, Porto, SSSA, and TUKL have carried out work. http://www.hurray.isep.ipp.pt/activities/art-wise/ rts.eit.uni-kl.de/research/mediaprocessing

Adaptive service configuration for Quality-of-Service aware middleware

Complex dynamic real-time scenarios may prevent the possibility of computing optimal service configurations before execution, an iterative refinement approach with the ability to trade off deliberation time for the quality of the solution was specified. The approach is to quickly find a good initial solution and to propose heuristic evaluation functions that optimize the rate at which the quality of the current solution improves as the algorithms have more time to run.

The work has also addressed the problem of dynamically changing system conditions, allowing the system to make QoS adaptation decisions in response to fluctuations in the nodes service load, under the control of the user. Monitoring the stability period and resource load variation of Service Level Agreements for different types of services is used to dynamically adapt future stability periods, according to a feedback control scheme. Work was done by Madrid and Proto (www.hurray.isep.ipp.pt/activities/qos)

Middleware

System adaptation requires full knowledge of the system state, therefore work has also been carried out in a framework to gather actual resource usage information, and interact with the operating system, extending the traditional POSIX trace model with a partial reflective model for operating system monitoring. The work was done by Porto.

HOLA-QoS is a framework for managing QoS and resources and it has been used in media processing which UPM and UC3m have developed jointly. It is implemented as a layered architecture, so that layers can be replaced, as far as the API is kept. The higher layers are meant to deal with quality, while lower layers are mainly related with resource management. In particular, the lowest layer is intended to manage budgets or resource shares assigned to applications. This layer has to provided accounting and enforcing facilities to ensure that budgets are guaranteed. Some times this functionality is provided by what is called resource kernels.

Cluster partners have developed kernels that provide these facilities and, hence, could be suitable to act as the lower layer of a HOLA-QoS based system. Work that it is under development is to port HOLA-QoS on top MARTE (Cantabria) and SHARK (Pisa) kernels. One result of this work is the possibility of experimenting with the adaptation techniques that these advanced resource kernels provide. Some publications on HOLA-QoS can be found at http://www.dit.upm.es/str.

Resource availability prediction

The resources typically used in-home entertainment applications (e.g., video/audio streaming) exhibit fluctuating availability. It is desirable to have mechanisms for indicating the available bandwidth during system runtime.

A comparative analysis of bandwidth estimation techniques for WiFi links has been carried out. In particular, the analyzed estimation techniques include several statistical and control-based algorithms. The analysis has identified the best suitable techniques taking into account the specific behavior of WiFi links. Work was carried out by UPC and TUKL. Analysis available at http://www.upcnet.es/~pmc16/nde_06.pdf.

2.2.2 Publications Resulting from these Achievements

Rosa Castañé, Pau Martí, Manel Velasco, Anton Cervin and Dan Henriksson, ”Resource Management for Control Tasks Based on the Transient Dynamics of Closed-Loop Systems”, ECRTS06, Dresden, Germany

G. Cecchetti and A. L. Ruscelli, “A Framework for Enhanced QoS Support in IEEE802.11e Networks”, Proc. of the 2nd International Mobile Multimedia Communications Conference (ACM-MobiMedia 2006), Alghero, Sardinia, Italy, September 2006.

Josep Guàrdia, Pau Martí, Manel Velasco and Rosa Castañé, “Enabling Feedback Scheduling in TrueTime”, ESAII-RR-06-15, Automatic Control Department, Technical University of Catalonia, Spain, March 2006

Rafael Camilo Lozoya Gámez, Pau Martí, Manel Velasco and Josep M. Fuertes, Wireless Network Delay Estimation for Time-Sensitive Applications, ESAII-RR-06-12, Automatic Control Department, Technical University of Catalonia, Spain, July 2006.
T Facchinetti, G. Buttazzo, L. Almeida. Dynamic Resource Reservation and Connectivity Tracking to Support Real-Time Communication among Mobile Units, EURASIP Journal on Wireless Communications and Networking, 2005(5): 712-730, December 2005.

Damir Isovic, Gerhard Fohler, Liesbeth Steffens, et. al., Temporal impact of high-level resources on end-to-end timing of video streaming, Technical Report, May, 2006

Pengpeng Ni, Damir Isovic, Gerhard Fohler: User-friendly H.264/AVC for Remote Browsing, ACM Multimedia 2006, Santa Barbara, CA, USA, October, 2006

R. Marau, L. Almeida, P. Pedreiras. Enhancing Real-Time Communication over COTS Ethernet switches. WFCS 2006, IEEE 6th Workshop on Factory Communication Systems, Turin, Italy. June 2006.

Luis Miguel Nogueira, Luis Miguel Pinho, "Iterative Refinement Approach for QoS-aware Service configuration", To be published in proceedings of the 5th IFIP Working Conference on Distributed and Parallel Embedded Systems (DIPES 2006), Braga, Portugal, October 2006.

Luis Miguel Nogueira, Luis Miguel Pinho, "Dynamic Adaptation of Stability Periods for Service Level Agreements", To be published in proceedings of the 12th IEEE International Conference on Embedded and Real-Time Computing and Applications (RTCSA 2006), Sydney, Australia, August 2006.

Larisa Rizvanovic, Gerhard Fohler, Real-time Architecture for Networked Multimedia Streaming systems, MiNEMA Workshop, 2006, Leuven, Belgium, February 2006

Filipe Valpereiro, Luis Miguel Pinho, "POSIX Trace Based Behavioural Reflection", Published in Proceedings of the 11th International Conference on Reliable Software Technologies (Ada-Europe 2006), Porto, Portugal, June 2006

2.2.3 Keynotes, Workshops, Tutorials

Keynote: The challenges of operational flexibility in real-time communication

Luis Almeida, University of Aveiro

CEDI 2005, Congreso Español de Informática
Granada, Espanha – 13-16th of September, 2005

This was the first congress of this kind in Spain, joining together a rather large set of different topics related to Informatics. The congress was organized like a federated set of mini-symposia or workshops running together, each dedicated to a specific topic, and with just a few global activities.

The talk was within the "Simposio de Sistemas de Tiempo-Real" and it addressed the issues related with improving operational flexibility in communication systems as required for distributed adaptation, from operational architectures to run-time mechanisms, timeliness guarantees, fault-tolerance, safety, dynamic QoS management and reconfiguration.

http://cedi2005.ugr.es/2005/programa_s20_str.shtml

Workshop: Requirements for Flexible Scheduling for Media Processing
Massy, France, June 16, 2006-09-08

A workshop on requirements for adaptive resource management was held together with the FRESCOR project. Part of the workshop focused on requirements for media processing**.**

http://soller.eit.uni-kl.de/mediawiki/index.php/Main_Page

3. Future Work and Evolution

3.1 Problem to be tackled over the next 18 months (Sept 2006 – Feb 2008)

In the next 18 months we will expand the application domain of the activity to more general media processing, to provide for more industrial input, including non mass market video processing and telecommunication.

We will continue to collect requirements to feed input to development of our adaptive methods, including the expanded application domain.

We will expand the integration of resources to be managed jointly. With respect to scheduling and cache, we will develop first algorithms to reflect on both scheduling and cache management. We will develop algorithms for adaptation of fluctuating resources, in particular wireless bandwidth and stream transformations.

The integration of HOLA-QoS with SHARK and MARTE will continue during this period of time, by UC3M and UPM. In the context of this activity, the aims will be to test the overall environment with multimedia applications to check the fulfilment of the quality properties. The used adaptation algorithms will be revised in order to take advantage of the resource management features provided by SHARK and MARTE.

UPM and UC3M will continue their work on extending Java-RMI with QoS management features. The goal is to improve remote invocations predictability with respect to selected QoS properties and to ensure that a server has enough resources to attend to a given number of invocations from a set of clients.

Further investigation will be done on resource reservation mechanisms, in order to extend them to other types of resources (not only processors) and to make them work under mutual exclusion constraints and the possibility of using the elastic scheduling theory in a energy-aware context, with the objective of balancing energy consumption with bandwidth requirements.

Work on the use of kernels developed by partners for HOLA-QoS will continue and focus in particular on the integration of communication aspects. Design of a component intended to adapt budgets at the low level API of HOLA-QoS to the functionality and model of the budgets of SHARK and MARTE. Implementation of this component.

Energy management: optimization of power consumption for CE devices is desirable. Adjustable CPU speed and memory speed can provide mechanisms for prolonging batteries life. The adaptive tuning of both speeds will be analyzed. UPC will carry out this work.

The work on anytime, iterative algorithms for QoS provisioning in collaborative embedded systems, and will address server-based scheduling approaches for collaborative service isolation will continue.

3.2 Current and Future Milestones

1. (achieved) Year1: Identify case studies, perform preliminary assessment
2. (achieved) Year2: Define a set of case studies and from them deduce the QoS requirements and their mapping into operational parameters of the computing and communication infrastructures.
   Temporal requirements for video streaming from stream demands and the temporal impact of devices have been identified. Adaptive methods for resource management and QoS middleware have been developed. Relations have been identified, e.g., on wireless networks.
3. **Year3: Expend these into a meaningful set of requirements of dynamic application domains (for instance, multimedia) that allow the creation of global mechanisms for resource management.**
4. Year4: Integration of the application adaptation processes into a general QoS resource management structure..

3.3 Indicators for Integration

The goal of this activity is to provide a uniform view on the requirements of some highly adaptive application environments, based on specific industrial needs.

Expected results and visible impact are: influence on integrating the operating system design, network, and middleware.

3.4 Main Funding

TUKL and Philips are partners in the STREP BETSY. http://cordis.europa.eu/fetch?ACTION=D&CALLER=PROJ_IST&QM_EP_RCN_A=71565

Currently, Aveiro has one national project that includes this topic (despite being substantially broader), CAMBADA – Cooperative Autonomous Mobile robots Based on an Adaptive Distributed Architecture, funded by the Portuguese governmental science agency: FCT. http://www.ieeta.pt/atri/cambada/

Since 2004, SSSA has one national project funded by the Italian Ministry of University Research (MIUR) on Real-Time Operating Systems for Supporting Cooperating Autonomous Robots, which also addresses resource reservation and adaptive bandwidth management policies.

FRESCOR IST project, in which the following ARTIST2 partners are involved: University of Cantabria, University of York, Scuola Superiore Santa Anna, Technical University of Kairserslautern, Technical University of Valencia and the Czech Technical University in Prague – www.frescor.org

THREAD Spanish project, in which the following ARTIST2 partners are involved: Technical University of Madrid, University of Cantabria, and Technical University of Valencia

3.5 Internal Reviewers

Giorgio Buttazzo

Pau Marti