ARTIST 2

Network of Excellence

IST-004527 ARTIST2:
Embedded Systems Design

Activity Progress Report for Year 2

JPRA-Cluster Integration

Flexible Scheduling Technologies

Clusters:

Adaptive Real-Time

Activity Leader:

Michael González Harbour (Universidad de Cantabria)
http://www.ctr.unican.es/people/michael.html

Many applications domains (including robotics, automotive and multimedia systems) require the execution of multiple concurrent activities with different criticality and timing constraints (e.g., periodic, aperiodic, time driven, event driven tasks) for which a single scheduling policy is not sufficient to satisfy the different requirements of the application.

Hence, the objective of the research is to exploit the excellence of the different teams for developing a real-time scheduling framework capable of handling various types of tasks with different real-time requirements in the same system. The challenge is to develop an efficient resource manager that can be adopted in next generation kernels to perform adaptive QoS control of time sensitive applications with dynamic characteristics.

Table of Contents

1. Overview of the Activity

1.1 ARTIST2 Participants: Expertise and Roles

Michael Gonzalez Harbour – University of Cantabria (Spain)
fixed-priority and hierarchical schedulability analysis, operating systems, and distributed systems

Gerhard Fohler – Kaiserslautern University (Germany)
combining online and offline scheduling and in resource management in consumer electronics

Giorgio Buttazzo – Sculoa Superiore Sant'Anna in Pisa (Italy)
dynamic scheduling schemes and management of quality of service requirements

Alan Burns – University of York (UK)
fixed priority schemes and real-time languages

Luis Almeida – University of Aveiro (Portugal)
network traffic scheduling

Eduardo Tovar – Polytechnic Institute of Porto (Portugal)
network scheduling and system wide analysis

1.2 Affiliated Participants: Expertise and Roles

Ivo De Lotto – University of Pavia (Italy)
dynamic priority schemes

Paolo Gai – Evidence (Italy)
real-time kernel for multiprocessor embedded systems, and in automotive systems

Lucia Lo Bello – Univ. of Catania (Italy)
communication protocols and stochastic scheduling

1.3 Starting Date, and Expected Ending Date

The activity started in September 2004 and it will be completed when the cluster is able to define a scheduling framework that is capable of managing quality of service requirements for multiple resources.

Expected ending date: March 2009

1.4 Baseline

Theory for independent scheduling algorithms is well defined in the areas of event triggered and time triggered systems, but few theoretical results have been achieved in trying to integrate these approaches. Some partial results exist for simplified architectures, but it is necessary to enhance them by taking into account all of the requirements of modern real-time systems including distributed ones. In addition to the development of theory, an experimental framework needs to be built in order to measure the performance of the different scheduling algorithms, and evaluate their applicability to real application domains.

1.5 Problem Tackled in Year2

Work has focused on the integration of diverse scheduling schemes, such as fixed priority based, dynamic priority based, and offline construction into a single coherent set that provides the advantages of the diverse approaches. Instead of creating an entirely new scheduling methodology, we primarily focus on methods to combine existing and widely applied schemes. Thus, the final results will transform the issue of scheduler selection and design from a single paradigm dominated selection or redesign to construction selecting and combining standard heterogeneous scheduling components.

After the work in year 1 consisting primarily on the development of new theory for the case of energy-aware scheduling, the work for the following years will focus on the integration of more resources into the flexible scheduling framework. Year 2 work has thus focused on the requirements that real-applications have in relation to this integrated framework, and in the selection of an appropriate base line for the scheduling framework.

1.6 Comments From Previous Review

1.6.1 Reviewers' Comments

The goal of this activity was to integrate various scheduling algorithms into a coherent set.

The activity progress report now clearly states the achievements and relationship between partners.

As mentioned in previous reports some techniques such as symmetric multiprocessing and virtualisation of operating systems might be considered in the future as potential for solving part of the issues addressed here, for instance power consumption reduction.

The deliverable is accepted as it is.

1.6.2 How These Have Been Addressed

The objectives of the work have been extended to handling multiple resources, including multiple processors, networks, memory, and energy.

2. Summary of Activity Progress

2.1 Previous Work

New scheduling mechanisms for integrating overload management techniques with energy-aware strategies were investigated in the context of real-time systems. The new scheduling mechanisms were analysed to guarantee timing constraints while minimizing energy consumption, and a kernel infrastructure was developed into the Shark operating system in order to facilitate their implementation. Moreover, the assessment of the (m,k)-firm model was done for its implementation over the FTT-Ethernet protocol.

2.2 Current Results

2.2.1 Technical Achievements / Outcomes / Difficulties encountered

Requirements for integrated-resource scheduling framework
A workshop on “Requirements for Flexible Scheduling in Complex Embedded Systems” Was held in Massy (Paris) in June 2006, with the objective of developing a set of requirements for building a flexible scheduling framework for applications demanding various types of tasks, constraints, and scheduling paradigms within the same system, and paying attention to the integration of multiple resources. See a description in the section on workshops, below. The workshop was very sucessful and brought together 20 participants from the following institutions:

- Technical University of Kaiserslautern, Germany
- Visual Tools, Spain
- University of York, UK
- Thales Communications France
- Evidence, Italy
- University of Cantabria, Spain
- Scuola Superiore Santa Anna Pisa, Italy
- Technical University of Madrid, Spain
- Polytechnic Institute of Porto, Portugal
- CEA, France
- Rapita Systems, UK
- University of Aveiro, Portugal
- Czech Technical University in Prague, Czcech Republic
- Technical University of Valencia, Spain

The results of the workshop, a wide collection of application requirements, are now in the process of being refined to produce a final document with a clear set of requirements for the integrated resource scheduling framework.

http://www.artist-embedded.org/FP6/ARTIST2Events/Events/flex_sched/

Baseline for integrated-resource scheduling framework
The FIRST IST project that finished in 2005 produced as its main result a contract-based scheduling framework that was capable of scheduling multiple application components with various kinds of requirements for CPUs and, to a limited extent, for networks in distributed systems. This framework has been selected as the baseline for the more ambitious framework

that will be developed in this activity and that will take into account the integrated scheduling of multiple resources. The FTT framework was also extended to micro-segmented switched Ethernet-based distributed systems, having revealed potential to provide efficient support to the contract model, to dynamic QoS management and to integrated resource scheduling in distributed environments. Moreover, the impact of flexible scheduling on dependability for distributed safety-critical applications was also assessed using FTT-CAN and appropriate mechanisms have been developed.

New theoretical developments

The contract-based scheduling framework defined above needs to be implemented using a specific scheduling strategy, and the most effective approach for this case is the server-based hierarchical scheduling in which an application or application component is scheduled over a protected bandwidth-preserving server (such as a periodic server, a sporadic server, or a constant bandwidth server) and individual threads in that component are scheduled by a higher-level scheduler that uses the bandwidth provided by the server. Theory has been developed towards being able to analize such scheduling schemes. In particular, methods have been developed in the University of York to analyze threads scheduled by fixed priority high-level schedulers based on different kinds of underlying fixed-priority schedulers. These methods have been extended by the University of Cantabria to EDF high-level schedulers based on constant-bandwidth servers. As a result, a complete analysis method exists for this kind of hierarchical scheduling.

Work has been done by the University of York together with Technische Universiteit Eindhoven (TU/e) on the analysis underpinning the use of CAN in real-time systems. CAN is now widely used is a range of real-time systems including high-integrity and flexible and adaptive systems. Unfortunately the analysis usually applied to its use has a fault that means that in some circumstances a system can be deemed schedulable when in fact it is not. This problem can arise because of the non-preemptive nature of the CAN protocol, and manifests itself when the first release of a message at a critical instant is not the worst case.

SSSA has developed the following theoretical results: energy-aware scheduling algorithms for processors with dynamic voltage scaling and discrete frequency levels; a method for minimizing the deadline of periodic tasks with the objective of reducing delay and jitter; a general methodology for performing sensitivity analysis of fixed priority periodic systems with configurable periods and computation times, allowing the system designer to derive the feasibility region of a task set and compute the maximum parameter variations that keep the system feasible.

Work carried out at the University of Aveiro has also exposed a couple of anomalies related to the definition of critical instant in hierarchical scheduling scopes found in communication systems that led to optimistic worst-case response time analysis in the past. Adequate methods were devised to cope with such anomalies. These results will be published in year3

Flexible architectures and communication protocols for networks used in distributed embedded real-time systems

This research has been done in collaboration between the following ARTIST 2 partners: University of Pavia, University of Catania (affiliated partner) and Malardalen University, Sweden. It consists of the following activities:

- Integration of networked subsystems in a resource constrained environment. This activity addressed the issue of achieving flexibility in networked subsystem integration through an integrated approach, where several systems are encapsulated as subsystems and later integrated on a shared hardware architecture. As the network is a resource shared by all

subsystems in the distributed architecture, its role in the integration process is particularly important, and the usage of an efficient and flexible network scheduler is essential. The core activity thus focused on assessing the suitability of Server-CAN, a network scheduler for the Controller Area Network, in the context of subsystem integration.

- Facilitating subsystem integration by decoupling priority and identifier in CAN messages. The CAN message identifier is especially important, as it not only does identify the message, but it also determines the message's priority. Hence, special care needs to be taken when assigning identifiers to messages. The research done on this topic resulted in a paper which outlines how CAN based systems are engineered today and indicates the potential and benefits of decoupling the message priority from the message identifier. Three solutions to this are given: TT-CAN, FTT-CAN and Server-CAN, and their strengths and weaknesses in an integration context are discussed.

2.2.2 Publications Resulting from these Achievements

R. I. Davis, A. Burns. "Hierarchical Fixed Priority Preemptive Scheduling", RTSS, pp. 389-398, 26th IEEE International Real-Time Systems Symposium (RTSS'05), December 2005

M. Aldea, G. Bernat, I. Broster, A. Burns, R. Dobrin, J. M. Drake, G. Fohler, P. Gai, M. González Harbour, G. Guidi, J.J. Gutiérrez, T. Lennvall, G. Lipari, J.M. Martínez, J.L. Medina, J.C. Palencia, M. Trimarchi. "FSF: A REAL-TIME SCHEDULING ARCHITECTURE FRAMEWORK". Proceedings of the 12th IEEE Real-Time and Embedded Technology and Applications Symposium, San Jose, CA, USA, April 2006.

Jose Luis Lorente and Jose Carlos Palencia. "An Edf Hierarchical Scheduling Model for Bandwidth Servers". 12th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications, Sidney, Australia, August 2006.

R.J. Bril, J.J. Lukkien, R.I. Davis, and A. Burns, "Message response time analysis for ideal controller area network (CAN) refuted", 5th International Workshop on Real Time Networks (RTN'06), Dresden, 2006.

Mauro Marinoni and Giorgio Buttazzo, "Adaptive DVS Management through Elastic Scheduling", Proceedings of the 10th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA 2005), Catania, Italy, September 2005.

Mauro Marinoni and Giorgio Buttazzo, "Balancing Energy vs. Performance in Processors with Discrete Voltage/Frequency Modes", Proceedings of the 12th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications, Sydney, Australia, August 2006.

Hoai Hoang, Giorgio Buttazzo, Magnus Jonsson, and Stefan Karlsson, "Computing the Minimum EDF Feasible Deadline in Periodic Systems", Proceedings of the 12th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications, Sydney, Australia, August 2006.

Enrico Bini, Marco Di Natale and Giorgio Buttazzo, "Sensitivity Analysis for Fixed-Priority Real-Time Systems", Proceedings of the 18th Euromicro Conference on Real-Time Systems (ECRTS 06), Dresden, Germany, July 2006.

T. Nolte, L. Lo Bello, H. Hansson, 'Integration of networked subsystems in a resource constrained environment', IEEE- IES the 11th IEEE International Conference on Emerging Technologies and Factory Automation ETFA 06, Sept. 20-22, 2006, Prague, Czech Republic.

T. Nolte, L. Lo Bello, H. Hansson, Facilitatinmg Subsystem integration by decoupling priorità and identifier in CAN messages, RTN 06, 5th Int. Worshop on Real-Time Networks, Satellite Workshop of ECRTS 06, July 4, 2006, Dresden, Germany.

Gerhard Fohler. Report on the results of the ARTIST2 Workshop on Requirements for Flexible Scheduling in Complex Embedded Systems, Massy (Paris), June 2006. http://www.artist-embedded.org/FP6/ARTIST2Events/Events/flex_sched/

R. Marau, P. Pedreiras, L. Almeida. Enhanced Ethernet Switching for Flexible Hard Real-Time Communication. RTN 2006, 5th Int. Workshop on Real-Time Networks (satellite to ECRTS'06), Dresden, Germany. July 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.

F. Carreiro, J. A. Fonseca, L. Almeida. Enhancing the Virtual Token Passing to Support Isochronous Traffic. WFCS 2006, IEEE 6th Workshop on Factory Communication Systems, Turin, Italy. June 2006.

Ferreira, J.; Almeida, L.; Fonseca, J.A.; Pedreiras, P.; Martins, E.; Rodriguez-Navas, G.; Rigo, J.; Proenza, J.. Combining Operational Flexibility and Dependability in FTT-CAN. IEEE Transactions on Industrial Informatics, 2(2):95-102, May 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.

L. Lo Bello, A. Gangemi. "A slot swapping protocol for time-critical internetworking", the Journal of Systems Architecture, Vol 51/9 pp 526-541, Elsevier, The Netherlands, Sept. 2005.

L. Lo Bello,G. Kaczynski, O. Mirabella, "Improving the Real-Time Behaviour of Ethernet Networks Using Traffic Smoothing", IEEE Transactions on Industrial Informatics, Special Section on Factory Communication Systems, Vol.1, NO.3, pp. 151-161, IEEE Industrial Electronics Society, USA, Aug. 2005.

M. Ryu, J. Park, S. Hong, L. Lo Bello, "Rapid Performance Re-engineering of Distributed Embedded Systems via Latency Analysis and K-Level Diagonal Search", Journal of Parallel and Distributed Computing, Volume 66, Issue 1, Pages 19-31, January 2006.

K.Kim, J. Diaz, L. Lo Bello, J. Lopez, C. Lee, S. Min, "An Exact Stochastic Analysis of Priority-Driven Periodic Real-Time Systems and Its Approximations ", IEEE Transactions on Computers, Vol.54, NO.11, pp.1460-1466, IEEE, USA, Nov 2005.

L. Lo Bello, O. Mirabella, "Scenarios and Architectures for Bluetooth Networks in Industrial Communication", IEEE ETFA '05, 10th IEEE International Conference on Emerging Technologies and Factory Automation, Catania, Italy, 19-22 Sept. 2005

T. Nolte, H. Hansson, L. Lo Bello, "Automotive Communications - Past, Current and Future", IEEE ETFA '05, 10th IEEE International Conference on Emerging Technologies and Factory Automation, Catania, Italy, 19-22 Sept. 2005.

L. Lo Bello, O. Mirabella, N. Torrisi," A General Approach to model traceability systems in food manufacturing chains", IEEE ETFA '05, 10th IEEE International Conference on Emerging Technologies and Factory Automation, Catania, Italy, 19-22 Sept. 2005

L. Lo Bello, O. Mirabella, A. Raucea, "ReTe: a Remote testbed for distance learning", IEEE ETFA '05, 10th IEEE International Conference on Emerging Technologies and Factory Automation, Catania, Italy, 19-22 Sept. 2005.

L. Lo Bello, F. Sgrò, G. Kaczynski, L. Di Stefano, O. Mirabella, "Improving Real-Time Behaviour of Ethernet switches using fuzzy Traffic Smoothing", Proceedings of the 2005 IFAC International Conference on Fieldbus Systems and their Applications, FET'05, Puebla, Mexico, Nov. 2005.

L. Lo Bello, O. Mirabella, “Clock Synchronization Issues in Bluetooth-Based Industrial Measurements”, Proceedings of the 6th International Workshop of Factory Communication Systems, IEEE WFCS 06, Turin, Italy, June 2006.

L. Lo Bello, G.A. Kaczynski, T. Nolte, “Towards a Robust Real-Time Wireless Link in a Land Monitoring Application”, in Proc. of ETFA06, the 11th IEEE International Conference on Emerging Technologies and Factory Automation ETFA 06, Prague, Czech Republic, Sept. 2006.

G. A. Kaczynski, L. Lo Bello, T. Nolte, “Towards Stochastic Response-Time of Hierarchically Scheduled Real-Time Tasks”, in Proc. of ETFA06, the 11th IEEE International Conference on Emerging Technologies and Factory Automation ETFA 06, Prague, Czech Republic, Sept. 2006.

L. Lo Bello, G. A. Kaczynski, F. Sgro', O. Mirabella, “A wireless traffic smoother for soft real-time communications over IEEE 802.11”, in Proc. of ETFA06, the 11th IEEE International Conference on Emerging Technologies and Factory Automation, Sept. 2006.

L. Lo Bello, O. Mirabella, “Efficient Full Duplex Links for Long Distance Wireless Mesh Networks”, IEEE- IES the 11th IEEE International Conference on Emerging Technologies and Factory Automation ETFA 06, Sept. 20-22, 2006, Prague, Czech Republic

2.2.3 Keynotes, Workshops, Tutorials

Keynote: Real-Time Issues in Mobile Wireless Networks
Speaker: Giorgio Buttazzo
Conference: 9th Int. Conference on Principles of Distributed Systems (OPODIS 2005)
Pisa, Italy - December 12-14, 2005

Keynote: Towards Component-Based Operating Systems
Speaker: Giorgio Buttazzo
Conference: Workshop on Operating Systems Platforms for Embedded Real-Time applications (OSPERT 2006)
Dresden, Germany - July 4, 2006

Keynote: Predictable Response Times in Event-Driven Real-Time Systems
Speaker: Michael González Harbour
Conference: Automotive – Safety & Security 2006
Stuttgart, Germany – October 2006
http://www.automotive2006.de/

Keynote : The challenges of operational flexibility in real-time communication
Speaker: Luis Almeida
Conference: Real-Time Systems Workshop at CEDI 2005, the 1st Spanish Congress on Informatics
Granada, Spain – September 15th, 2005
http://cedi2005.ugr.es/2005/invitados.shtml

Workshop : Requirements for Flexible Scheduling in Complex Embedded Systems
Massy(Paris), France – June 2006

As the complexity of embedded applications evolves, the gap between their timing requirements and the scheduling services provided by real-time operating systems increases. Today's systems require a mixture of quality of service, hard and soft real-time, and multiple resource management, which bring the need for raising the level of abstraction of the scheduling services. The objective of this workshop was to develop a set of requirements for building a flexible scheduling framework for applications demanding various types of tasks, constraints, and scheduling paradigms within the same system. The framework is oriented towards the contract-based scheduling of multiple resources such as the processors, the networks, dynamically reconfigurable modules, interrupts with time protection, shared resources with time protection, memory protection, and energy/power-aware scheduling.The framework should ensure that the techniques developed can be used in different application domains. The targeted application domains are at least industrial control systems, media processing applications, automotive embedded systems, telecommunications, and artificial intelligence. The set of requirements developed in the workshop should allow ongoing and future research and development projects to focus on the real needs of embedded systems, collected through the consensus of a large group of interests and experts, brought together through the ARTIST2 Network.
http://www.artist-embedded.org/FP6/ARTIST2Events/Events/flex_sched/

3. Future Work and Evolution

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

After the collection of the requirements for the flexible scheduling framework for integrated resources, the first step in the work over the next period is the refinement of the requirements to come out with a crisp set of concrete requirements that provide the widest possible coverage of the needs described, and that is implementable with moderate effort and reasonable oberhead in a scheduling framework on a particular kernel.

Once the final set of requirements has been established, it is necessary to develop an architectural model of the framework and the applications, and describe the interfaces among them. The approach that is considered more effective is based on the contract-based scheduling framework developed by some of the partners in the past FIRST EU project. This framework provides the ability to compose several applications or components into the system, and to flexibly schedule the available resources while guaranteeing hard real-time requirements. 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.

However, the FIRST framework represents just an initial step because it only considers the scheduling of two resources: CPU and network (the latter, to a limited extent). Therefore, one of the goals of the work in this period is to extend the contract based framework to provide a contract model that integrates other resources. The new framework will include all system resources in addition to the threads and networks: dynamically reconfigurable modules, multiple processors, interrupts with time protection, shared resources with time protection, memory protection, and energy/power-aware scheduling. Of course it is necessary to distinguish those resources that we can manage for scheduling purpose from those that we can just model, or that do not require specific scheduling.

The third goal in this period is to start the design of a quality of service manager that understands about the quality concepts of the application and translates them to the scheduling domain information that the underlying system can understand and implement, generating the contracts in an automatic or semiautomatic way. Of course application level quality varies from one application to another so this quality manager has necessarily a generic interface, and needs specific instantiation for a particular application. The manager provides a common reference architecture in which the core operations can be shared.

There is a strong relation between the flexible scheduling activity in ARTIST2 and the new FRESCOR EU project that started june 2006. Many of the partners in the activity are also partners of the project, and the goals are somehow shared. Indeed ARTIST2 has contributed to creating the necessary critical mass to set up a project like FRESCOR. Given the shared goals it is necessary to define clearly the role of each project in the overall activity. We see the following role of the flexible scheduling activity in ARTIST2:

- Bring together a wide body of expertise from the whole ARTIST2 community to come out with a broad and ambitious set of requirements for the flexible scheduling framework
- Use that expertise to establish and later evaluate the usefulness and applicability of a contract-based specification for the framework, that allows applications to express their requirements, and the scheduler to satisfy them in the most optimal way.

- Use the ARTIST2 community as one of the ways of disseminating the results of the FRESCOR project

With these roles in mind we think that this is an ideal situation in which the FRESCOR project can benefit from the ARTIST2 NoE expertise, and the network can benefit from being able to influence the project, and from being able to exploit its results.

3.2 Current and Future Milestones

- Year1: Preliminary work on the integration of diverse scheduling schemes. Achieved for the case of CPU and energy, and CPU and network
- Year2: Demonstrate the combination of specific scheduling schemes applied both to CPU as well as to the network, to suit diverse application requirements in the same system. This has been achieved both in theoretical developments for server-based hierarchical scheduling, and in practice through the FIRST scheduling framework.
- **Year 3: Define a set of requirements and an architectural model for the framework for flexible scheduling that integrates multiple resources, including CPUs and networks: dynamically reconfigurable modules, multiple processors, interrupts with time protection, shared resources with time protection, memory protection, and energy/power-aware scheduling.**
- Year4: Provide a framework that allows the seamless integration of flexible scheduling schemes for integrated resources, allowing the choice of appropriate scheduling methods for individual activities in the different resources.

3.3 Indicators for Integration

Joint results on scheduling algorithms, analysis tools, and kernel and network support for the integration and coexistence of diverse system-wide scheduling schemes.

3.4 Main Funding

In addition to the specific funds for the ARTIST2 NoE, the main sources of funding are coming from:

- 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
- THREAD Spanish project, in which the following ARTIST2 partners are involved: Technical University of Madrid, University of Cantabria,Technical University of Valencia
- DIRC (Dependability Intersisciplinary Reasearch Collaboration), UK project involving ARTIST2 partner University of York
- National project PRIN-MIUR 2004: "Real-Time Operating Systems for Supporting Cooperating Autonomous Robots". Project No. 2004095094_003, involving involving ARTIST2 partner SSSA.