ARTIST 2

Network of Excellence

IST-004527 ARTIST2:
Embedded Systems Design

Activity Progress Report for Year 1

JPRA-Cluster Integration:

Flexible Scheduling Technologies

Cluster:

Adaptive Real-Time

Activity Leader:

Giorgio Buttazzo (University of Pavia)

Many applications domains (including robotics, automotive and multimedia systems) require the execution of many 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. Introduction

1.1 Activity Leader

Team Leader: Giorgio Buttazzo (Pavia)
Areas of his team's expertise: dynamic scheduling and overload management techniques.

1.2 Policy Objective

Many applications domains (including robotics, automotive and multimedia systems) require the execution of many 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.

1.3 Industrial Sectors

Consumer electronics will directly benefit in the near future of the results of this research for improving the functionality of multimedia applications. The automotive industry is also interested in this research to handle overload conditions that frequently occur in the microcontrollers embedded in the car. Finally, flexible scheduling techniques are also demanded in industrial automation, where often robotics applications consists of several tasks with different criticality and timing constraints.

2. Overview of the Activity

2.1 Artist Participants and roles

Team Leader: Gerhard Fohler (Mälardalen University)
Areas of his team's expertise: offline scheduling.

Team Leader: Giorgio Buttazzo (Pavia)
Areas of his team's expertise: dynamic scheduling schemes.

Team Leader: Alan Burns (York)
Areas of his team's expertise: fixed priority schemes.

Team Leader: Michael Gonzalez Harbour (University of Cantabria)
Areas of his team's expertise: fixed priority analysis.
.
Team Leader: Luis Almeida (University of Aveiro)
Areas of his team's expertise: network traffic scheduling.

Team Leader: Eduardo Tovar (Polytechnic Institute of Porto)
Areas of his team's expertise: network scheduling and system wide analysis.

2.2 Affiliated partners and Roles

Team Leader: Giuseppe Lipari (Scuola Superiore S. Anna)
Areas of his team's expertise: dynamic priority schemes.

Team Leader: Paolo Gai (Evidence)
Areas of his team's expertise: real-time kernel for embedded systems.

Team Leader: Laurent Pautet (Paris Telecom)
Areas of his team's expertise: distributed systems.

Team Leader: Lucia Lo Bello (Univ. of Catania)
Areas of his team's expertise: communication protocol and stochastic scheduling.

2.3 Starting date, and expected ending date

It started in September 2004 and it will be complete when the cluster will be able to define a scheduling framework that integrates off-line, fixed-priority and dynamic-priority scheduling schemes in a single kernel.

Expected ending date: March 2007.

2.4 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.

2.5 Technical Description

Work will focus 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 entirely new scheduling methodology, we will primarily focus on methods to combine existing and widely applied schemes. Thus, the 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.

3. Activity Progress Report

3.1 Work achieved in the first 6 months

In the first 6 months, the ART cluster has contributed to progress in the following areas:

1. Energy-aware strategies to guarantee timing constraints while minimizing energy consumption. New scheduling mechanisms for integrating overload management techniques with energy-aware strategies have been investigated in the context of real-time systems. This result was achieved thanks to the collaborative work of Pavia, Pisa (affiliated to Pavia), and Lund. In particular, the added value of the ARTIST2 network was to combine the expertize on feedback control available in Lund and the expertize on operating systems available in Pavia and Pisa to increase system adptivity by integrating feedback techniques at the operating system level.
2. Resource reservation mechanisms to reduce intertask interference and provide temporal protection. The Constant Bandwidth Server has been investigated for a possible use in control applications to reduce the effects of overruns in tasks with variable computation times. Such methods have also been considered for handling media processing. This result was achieved by the joint work of Pavia, Pisa (affiliated to Pavia), Malardalen, and Cantabria, and Aveiro. Although each individual group had expertize in scheduling and resource management, the ARTIST2 network enabled these groups to concentrate their efforts on a larger and more significant problem, providing solutions directly usable in the domain of multimedia systems and consumer electronics.
3. Tools for schedulability analysis of complex systems. The main characteristics of a powerful schedulability analyzer have been discussed to provide sensitivity analysis during the design of complex real-time systems. This result was achieved by the joint work of Malardalen, Evidence (affiliated to Pavia), and York. Achieving this result, required the expertize in off-line scheduling (available in Malardalen), execution time analysis (available in York), and schedulability analysis tools (available in Evidence).

3.2 Work achieved in months 6-12

In months 6-12, the ART cluster investigated the following issues:

1. New scheduling mechanisms have been analysed to guarantee timing constraints while minimizing energy consumption. This result was achieved by integrating the work of Pavia on power-aware scheduling, of Malardalen on flexible timing constraints, of Pisa (affiliated to Pavia) on operating systems, and of Cantabria and York on schedulability analysis.
2. A kernel infrastructure has been developed into the Shark operating system in order to facilitate the implementation of energy-aware scheduling. This result was achieved by the joint work of Pavia, Pisa (affiliated to Pavia), and Evidence (affiliated to Pavia), all expert in kernel implementation.
3. The assessment of the (m,k)-firm model was done for its implementation over the FTT-Ethernet protocol. This result was achieved by the joint work of Pavia, Aveiro, Porto, and Catania (affiliated to Pavia). Here, the ARTIST2 network was essential for integrating the expertize of Aveiro and Porto on Ethernet protocols, of Catania on distributed systems analysis, and Pavia on overload management.

4. A contract-based framework for flexible scheduling has been discussed to be used in real-time distributed systems. This result was achieved by the joint work of Cantabria, Malardalen, Aveiro, York, and Pisa (affiliated to Pavia), who organized a meeting in Palma de Mallorca to deeply address such a new approach. Such a constructive discussion and brainstorming could not be done remotely in an efficinet way, hence the derived conclusions could not be achieved without the support of ARTIST2.

3.3 Difficulties Encountered

The major problem was encountered in measuring the benefits of the proposed algorithms in terms of power consumption, due to the difficulty of isolating the power consumption of the processor from the power consumption of the entire computer.

3.4 Recommendations

The work will proceed to extend the proposed kernel infrastructure to include other devices other than the CPU, as I/O devices, memory, actuators, and wireless communication boards.

3.5 Milestones

The major technical result was the development of a kernel infrastructure capable of supporting new scheduling algorithms for overload and energy management.

3.6 Main Funding

Main sources of funding are coming from two national projects carried out with other Italian Universities.

3.7 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.

3.8 Evolution

The long term perspective is to provide a framework providing for the seamless integration of flexible scheduling schemes, allowing the choice of appropriate scheduling methods for individual activities in a system or messages on the network.

The 18 month perspective is to 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 activity is expected to last beyond the period of initial funding for integration.

4. Detailed Technical View

4.1 Brief State of the Art

The main reason for investigating flexible scheduling technologies in a real-time embedded system is to provide predictability and adaptivity for systems and environments where requirements on resources are inherently unstable and difficult to predict in advance. Such a difficulty is due to different causes. First of all, modern computer architectures include several low-level mechanisms that are designed to enhance the average performance of applications, but unfortunately introduce high variations on tasks' execution times. In other situations, as in multimedia systems, processes can have highly variable execution times. As a consequence, the overall workload of a computing system is subject to significant variations, which can produce an overload and degrade the performance of the entire system in an unpredictable fashion. This situation is particularly critical for small embedded devices used in consumer electronics, telecommunication systems, industrial automation, and automotive systems. In fact, in order to satisfy a set of constraints related to weight, space, and energy consumption, these systems are typically built using small microprocessors with low processing power and limited resources.

For most of these systems, the classical real-time approach based on a rigid off-line design, worst-case assumptions and a priori guarantee would keep resources unused for most of the time, therefore is not acceptable for efficiency reasons. When resources are scarce, they cannot be wasted. On the other hand, an off-line design based on average-case behavior is also critical, because it would be difficult to guarantee timing constraints when resources are overloaded.

To prevent unpredictable performance degradations due to overloads, a real-time system must react to load variations, degrading its performance in a controlled fashion acting on system, as well as application parameters. The process of controlling the performance of a system as a function of workload variations is referred to as Quality of Service (QoS) Management. Performing efficient QoS management requires specific support at different levels of the system architecture. Hence, new software methodologies are emerging in Embedded Systems, which strictly relates to Real-Time Operating Systems (RTOS), Middleware, and Networks.

Real-time scheduling is the kernel mechanism having the most impact on RTOS performance. Most scheduling algorithms have been developed around one of three basic schemes: table driven, fixed priority, or dynamic priority. Depending on whether scheduling decisions are resolved before or during runtime, they are classified as offline, or online.

Table driven methods are capable of managing distributed applications with complex constraints, such as precedence, jitter, and end-to-end deadlines. As only a table lookup is necessary to execute the schedule, process dispatching is very simple and does not introduce large runtime overhead. On the other hand, the a priori knowledge about all system activities and events may be hard or impossible to obtain. Its rigidity enables deterministic behavior, but limits flexibility drastically. This approach is the one usually associated with a Time-Triggered architecture, such as TTP, which is commercially available.

Online methods overcome these shortcomings and provide flexibility for partially or non-specified activities. A large number of schemes have been described in the literature. These scheduling methods can efficiently reclaim any spare time coming from early completions and allow handling overload situations according to actual workload conditions. They can be distinguished into two main classes: fixed-priority and dynamic-priority algorithms. Fixed priority scheduling (FPS) is similar to many standard operating systems, assigning before runtime of the system priorities to tasks, and executing the task with the highest priority to execute from the set of ready tasks at runtime. Dynamic priority scheduling typically selects the ready task with the closest deadline; priorities do not follow a fixed patterns, but change dynamically at runtime. To keep feasibility analysis computationally tractable, however, tasks cannot have arbitrary constraints.

Adapting to changing environmental situations may involve changes to task parameters at runtime. System wide changes, e.g., for changing operational modes in the system, have been addressed by mode change algorithms.

Feedback scheduling changes task parameters, in particular periods, to respond to online variations in the environment and current load conditions of the system. As both conditions can vary frequently, too frequent responses, which in turn influence the conditions, can introduce instability in the system. Feedback control scheduling applies control theory to estimate effects of changes and to choose parameters to provide for smooth responses and avoid instability.

Each of the basic scheduling paradigms has specific advantages. When advantages of different schemes are demanded in the same system, more than one scheme could be used for different tasks. For example, in a complex system including hard periodic and soft aperiodic tasks, two scheduling schemes need to be integrated for satisfying the different requirements of each task class.

In hierarchical scheduling a meta algorithm arbitrates between a set of diverse scheduling algorithms. Thus, it can appear to the individual scheduling algorithms and their applications that they execute alone in the system. Furthermore, the amount of the CPU portion can be set individually for each scheduler and application. Special attention has to be given to shared resources.

Flexible scheduling is an underlying theme to most novel scheduling trends which goes beyond the standard model of completely known tasks with timing constraints expressed as periods and deadline, providing “yes or no” answers whether all tasks meet their deadlines. Issues addressed include probabilistic parameters, handling of applications with only partially known properties, relaxed constraints and such that cannot be expressed solely by periods and deadlines, coexistence of activities with diverse properties and demands in a system, combinations of scheduling schemes, and adaptation or changes to scheduling schemes used at run-time.

4.2 Industrial Needs and Experience

Needs for flexible scheduling techniques are typical of industries working in consumer electronics, industrial automation, and telecommunications, as resulted from a study carried out within the ARTIST 5FP project.

For example, mobile terminals today are getting more and more advanced, and features are added at an unprecedented speed. The code base is thus constantly increasing and a 3G feature phone of today typically consists of several million lines of code with use-cases involving large number of parallel activities. Getting all this to work is the delicate task of a few highly competent craftsmen, and it is not getting any easier. For these applications, the use of flexible real-time scheduling techniologies would allow to safely mix real-time and non real-time processes. The benefit of such a solution would be a much more scalable platform. Adding and removing features would become predictable and less hazardous, allowing configuring the system without worrying about unpleasant surprises.

In the area of Industrial Automation, the continuous increment in processing power and memory capacity in local processors gives the opportunity to add new tasks into them, increasing system complexity in terms of supervision, diagnostics, presentation, communication, etc. Adaptive tasks scheduling that preserves the real-time constraints is a possible way to handle such situation and to give insight about how manage the complexity of the tight execution of the different tasks in the real-time processing. Adaptive real-time systems technologies are very much used in many applications in the area of Industrial Automation.

In telecommunication companies, embedded systems are mainly targeted to the interfaces between communication technologies or media and to coding/decoding operations. Most of them work at the telecommunications company side, although some are also installed in customer facilities. They may be considered real-time as they have timeliness requirements for some of the critical operations they must perform. Although it is not a key issue at the management level of the companies, real-time operation is a concern for telecommunications embedded systems designers. The main current interest seems to be in exploring the use of real-time extensions for the Linux OS. It also seems that QoS mechanisms are starting to be recognized as important for these embedded applications, namely in order to increase the efficiency of subsystems and to support the possibility to serve more clients with similar levels of resources.

4.3 Ongoing Work in the Partner Institutions

At the Univeristy of Pavia, the work on flexible scheduling is focusing on new methodologies for integrating overload management techniques with energy-aware strategies, in the context of small embedded systems for battery operated devices. New scheduling mechanisms have been proposed and analysed to guarantee timing constraints while minimizing energy consumption, and a kernel infrastructure has been developed into the Shark operating system in order to facilitate their implementation. The work will proceed to extend the proposed kernel infrastructure to include other devices other than the CPU, such as I/O devices, memory, actuators, and wireless communication boards.

The Scuola Superiore Sant'Anna (affiliated to Univeristy of Pavia) is active on resource reservation mechanisms that reduce intertask interference and provide temporal protection among the concurrent activities. Such a mechanism is essential to handle execution overruns in applications where tasks have higly variable computation times.

Evidence s.r.l. (affiliated to Univeristy of Pavia) is providing support for implementing such novel scheduling techniques on kernels that have to run on small microprocessors with scarce resources. In addition, Evidence is developing tools for simplifying the programming and the analysis of such complex systems.

MDH continued to work on adaptive resource managment for media processing. One of the core topic was applying real-time methods for adaptation at the nework layer. MDH was active in the EU IST project FIRST, which was completed during the period. In this time, MDH worked on integrating offline and online scheduling, and in particular on the coexistence of fixed priority and table-driven scheduling.

The University of Cantabria is been working on a contract-based framework for flexible scheduling in real-time networks. The framework provides the ability to make dynamic bandwidth reservations using a distributed acceptance test that ensures the overall network schedulability. A prototype implementation has been implemented using the Real-Time

Ethernet Protocol (RT-EP), which is a software-based token-passing Ethernet protocol for real-time applications, that does not require any modification to existing Ethernet hardware. The implementation includes a distributed negotiation mechanism as well as a new scheduler that manages budget consumption and enforces the bandwidth reservations. The implementation has proved to be feasible, efficient, and easy to use. It has been successfully used to implement a distributed controller for an industrial robot.

The team at the University of Aveiro has been researching in two different lines. On one hand, the team has addressed the image processing associated to robotic vision and control, a situation that is known for its varying load. An approach based on appropriate process partitioning and scheduling has been developed and deloyed in the RoboCup Middle-size team of this University. On-going work addresses the online adaptation of the processes attributes to enhance the quality-of-service of the robot behavior that is active in each instant. On the other hand, the team is developing a dynamic adaptation scheme for control streams of distributed embedded control systems that reduces the bandwidth requirements while maintaining the control performance.

4.4 Interaction, Building Excellence Between Partners

The interaction among the ART partners within Artist2 is being extremely helpful for building a general view of the existing scheduling panorama, with the advantage of perceiving the approaches followed by the partners under different perspectives. This generates interesting discussions that are extremely constructive in the process of integrating several algorithms together in the same system. In same case, such interactions generate new ideas on applying existing algorithms to different domains (e.g., from processor to networks, or viceversa) or extending the algorithms to more general scenarios.

4.5 Spreading Excellence

Excellence will be spread both in the industrial and academic environment. Several partners in the ART cluster have already tight contacts with the major consumer electronics industries (like Philips and Ericsson) and plan to organize meetings and workshops to transfer the acquired know how.

In the academic world, such novel methodologies will be spread in university courses on real-time systems (already tought by most of the partners in their respective sites) and by organizing summer schools on specific topics for graduate and undergraduate students.