Journal of Embedded Computing 1 (2005) 167–168
IOS Press

Guest-editorial

Special issue: Real-Time Systems (Euromicro RTS-03)

Giorgio Buttazzo
Guest-editor

Embedded computing systems are becoming pervasive in our society, spreading in several novel application domains, including multimedia systems, consumer electronics, advanced simulations, automotive systems, sensor networks, virtual reality, intelligent toys, and entertainment robotics.

In most applications, the computing system is required to work under stringent resource constraints, imposed by the characteristics of the specific system. For example, battery operated devices must limit energy consumption for ensuring a long lifetime; the equipment embedded in small robots must have a light weight for facilitating mobility; high quality digital televisions or video-projectors must limit dissipation to avoid noisy fans.

In spite of such constraints, most embedded computers have to interact with the controlled system through a number of sensors and actuators, hence they are inherently real time and are required to meet timing constraints imposed by the environment.

Providing high responsiveness under stringent timing and resource constraints is the main challenge in today's embedded systems, which requires the joined contribution of different research areas, as real-time scheduling, control theory, and distributed systems.

The papers collected in this special issue provide an important support for addressing the problems outlined above. They have been selected among the best papers presented at the Fifteenth Euromicro Conference on Real-Time Systems, the premier European conference in the area of real-time computing, held in Porto, Portugal, between July 2 and 4, 2003. The versions presented here have been revised and extended by the authors to improve the quality of presentation and include the most recent results of the on going research.

The first paper, ”Multiprocessor fixed-priority scheduling with restricted interprocessor migrations”, by Sanjoy Baruah and John Carpenter, considers the problem of scheduling periodic and sporadic tasks upon identical multiprocessors, under the assumptions that each task instance can only have a fixed priority and can only execute on a single processor. These restrictions can be used when the overhead due to priority computation and task migration needs to be reduced.

The second paper, “Probabilistic Timing Analysis: an Approach Using Copulas”, by Guillem Bernat, Alan Burns, and Martin Newby, proposes a statistical tool (copulas) for analyzing the stochastic behavior of complex real-time systems characterized by variable execution times. Copulas are used for determining the probability distribution of the worst-case execution time of a real-time program and allow the description of the dependence structure between blocks.

The third paper, “The Capacity of an Implicit Prioritized Access Protocol in Wireless Sensor Networks”, by Marco Caccamo and Lynn Zhang, enables reliable monitoring and intelligent control of the physical environment through a distributed network of wireless sensors/actuators. In particular, this paper focuses on the problem of providing delay and throughput guarantee to real-time messages and analyzes the sensor network capacity when messages are scheduled with the Implicit-EDF algorithm.

The fourth paper, “Control-Scheduling Codesign of Real-Time Systems: The Control Server Approach”, by Anton Cervin and Johan Eker, combines real-time scheduling and control theory for improving the per-

ISSN 1740-4460/05/$17.00

formance of complex control applications. The paper presents a real-time scheduling mechanism (the control server) which creates the abstraction of a control task with a specified period and a fixed input-output latency shorter than the period. Individual tasks can be combined into more complex components without loss of their individual guaranteed fixed-latency properties.

The fifth paper, “Response Time Analysis of EDF Distributed Real-Time Systems”, by J.C. Palencia and M. Gonzlez Harbour, presents a general methodology to verify the feasibility of the schedule in a periodic set of tasks characterized by release offsets. This technique is extended to dynamic priority scheduling and can be used to perform a global schedulability analysis of different task models in distributed heterogeneous systems, where processors and communication networks may use fixed or dynamic priority.

The sixth paper, “Real-time issues of MPEG-2 playout in resource constrained systems”, by Damir Isović, Gerhard Fohler, and Liesbeth Steffens, presents results from a study of realistic MPEG-2 video streams to analyze the validity of assumptions typically made for software decoding. Moreover, the paper identifies a number of misconceptions and constraints imposed by frame buffer handling and discusses their implications on decoding architecture and timing.

The seventh paper, “A methodology for designing hierarchical scheduling systems”, by Giuseppe Lipari and Enrico Bini, addresses the problem of composing different real-time applications on a single processor system while guaranteeing that their timing requirements are not violated. The authors adopt a resource reservation approach, where each application is handled by a dedicated server that is assigned a fraction of the processor.

The eighth paper, “Multi-version Scheduling in Rechargeable Energy-aware Real-time Systems”, by Cosmin Rusu, Rami Melhem, and Daniel Mossé, deals with battery-powered real-time systems, and proposes a method for balancing energy, deadlines and task rewards to achieve a variety of QoS-aware tradeoffs. Assuming that energy can be replenished by using a rechargeable battery, the authors propose both a static solution that maximizes the system value in a worst-case scenario, and a dynamic scheme that takes advantage of the extra energy in the system when worst-case scenarios do not happen.

The ninth paper, “Efficient Scheduling of Soft Real-time Applications on Multiprocessors”, by Anand Srinivasan and James Anderson, considers soft real-time applications implemented on multiprocessors. Since Pfair scheduling algorithms are the only known means of optimally scheduling *hard* real-time applications on multiprocessors, the paper focuses in efficient implementations of Pfair scheduling algorithms, which trade optimality with acceptable runtime overhead.

A final acknowledgement goes to the anonymous reviewers, who provided valuable suggestions to improve the quality of the presented material.

The Guest-editor,
Giorgio Buttazzo
University of Pavia, Italy
E-mail: buttazzo@unipv.it