Guest editorial

Daniel Mossé · Julius C.B. Leite · Dara Kusic

Published online: 19 February 2011
© Springer Science+Business Media, LLC 2011

In the past decade, energy-awareness has become a first-class requirement of efficient real-time systems. Real-time researchers have contributed significantly to the development of such systems, with particular focus on maintaining real-time system properties including hard and soft task deadlines, reliability, and time-based service-level objectives. Power awareness and energy efficiency are not only beneficial for increasing the lifetime of embedded systems, but they can also help the handling and energy-efficient scheduling of devices attached to such systems. This special issue on “Energy-Aware Real-time Systems” presents the state-of-the-art in the area of energy and power management in real-time computing systems.

This collection of papers develops instrumental techniques for designing and implementing energy-efficient real-time systems. As the current generation of transistors nears the atomic limits of gate size, the role of software designers becomes evermore integral to managing the limited power supply available to computing devices. The goal is to develop and assemble an array of energy-conserving techniques aimed at utilizing every working cycle of a component’s time toward completing a task; else, the component is sent to a low-power or no-power state until it is needed. The best

At the time the work was commenced, D. Kusic was supported by an NSF/CRA Computing Innovations Fellow (CIFellow) grant number 0937060.

D. Mossé · D. Kusic (✉)
University of Pittsburgh, Pittsburgh, USA
e-mail: dmk64@pitt.edu

D. Mossé
e-mail: mosse@cs.pitt.edu

J.C.B. Leite
Universidade Federal Fluminense, Niterói, Brazil
e-mail: julius@ic.uff.br

of these approaches consider selecting the right energy-management techniques for the current workload conditions, developing formulaic abstractions for energy management that can apply to a broad family of electronic components, and accounting for issues of low-power electronic reliability by anticipating task time for redundant computations.

Among the many papers submitted with valuable contributions to make in energy-aware system design, this special issue could only accommodate a small number of works that cover energy and power management approaches in a range of system architectures. The four selected papers present important recent results that address energy-efficient design in a range of real-time system scenarios, applying a variety of enabling technologies.

- The first paper, “Hybrid Power Management in Real Time Embedded Systems: An Interplay of DVFS and DPM Techniques,” applies to embedded systems a machine-learning technique at run time to dynamically select the best power management policy for a given workload, while still maintaining the specified timing requirements.
- The second paper, “Energy Efficient Scheduling for Real-Time Embedded Systems with QoS Guarantee,” deals with soft real-time task deadlines, applying a combination of power management techniques to turn processors on/off while taking into account the overheads of applying power management techniques, and prioritizing which task instances to drop when the system is overwhelmed.
- The third paper, “Applying Real-Time Interface and Calculus for Dynamic Power Management in Hard Real-Time Systems,” also deals with changing the power state of devices, but based on a flexible real-time calculus model that has wide applicability to deadline constrained systems.
- The last paper, “Global Scheduling Based Reliability-Aware Power Management for Multiprocessor Real-Time Systems,” adds quality guarantees to power-aware and energy-aware real-time systems by introducing reliability allowances within the scheduling algorithm, so that the unused slack can be used either to save energy or to provide fault tolerance in the form of redundant computations.

In addition to these important developments, research on real-time energy and power mechanisms in the future will spread to the multi-core systems with their shared caches, to memory subsystems that are consuming increasingly more overall energy (due to their sizes and to the advances in CPU energy and power management), and also to clusters of computers that are rising in importance (due to the popularity of cloud computing and virtual-machine-based systems).

Finally, we would like to thank the authors of all submitted papers and express our gratitude to the referees for their thorough review work. This special issue would not have been possible without the support of the editorial review board and staff of the Real-Time Systems journal.