Real-Time Syst (2016) 52:563–565
DOI 10.1007/s11241-016-9249-0

Special issue of the Euromicro Conference on Real-Time Systems (ECRTS)

Rolf Ernst[1]

Published online: 8 February 2016

The Euromicro Conference on Real-Time Systems (ECRTS) is a premier forum for research in real-time systems. The 26th Edition was held in Madrid 2014 from July, 8–11. Out of the 114 submissions with authors from 25 countries, the program committee selected 27 papers for presentation. Out of these papers, 4 top rated papers were selected and invited to submit an extended paper to this special Issue. These invited papers, then, were carefully peer reviewed before they were accepted as printed here.

The paper,

> "Quasi-Partitioned Scheduling: Optimality and Adaptation in Multiprocessor Real-Time"
> by Ernesto de Souza Massa, George Lima, Paul Regnier, Greg Levin, and Scott Brandt

presents an adaptive multiprocessor scheduler that exploits the variable load of sporadic task sets to reduce the run-time overhead. In high load situations, where tasks arrive with their minimum interarrival times, a global scheduling strategy using Quasi-Partitioned Scheduling (QPS) servers for predefined task sets is used. For increasing interarrival times, the QPS servers are deactivated and switch to a partitioned scheduling with local EDF. The extensive evaluation with synthetic task sets shows that this scheduling strategy provides a smooth reduction of migration and preemption overhead as the interarrival times increase and the load goes down. This effect is observed from small to large processor configurations. Such an adaptive combination of global and local strategies is new and opens a new dimension for research.

✉ Rolf Ernst
r.ernst@tu-bs.de

[1] Technische Universität Braunschweig, Braunschweig, Germany

Partitioned caches have been proposed to make preemptive multi-tasking systems predictable as compared to systems with unconstrained caches which are hard to predict. However, part of the advantage of partitioned caches might just be due to the harder WCET analysis problem of non-partitioned caches. The paper,

"On the Effectiveness of Cache Partitioning in Hard Real-Time Systems"
by Sebastian Altmeyer, Roeland Douma, Will Lunniss, and Rob Davis

employs a recent WCET analysis tool for unconstrained caches and compares the performance with a partitioned cache for which they develop a tool for optimal cache partitioning. The metrics they use is schedulability of task sets. Extensive experiments with different benchmark sets and synthetic examples demonstrate that in most cases the unconstrained caches provide better worst-case schedulability, indicating that the advantage of an increased predictability does not compensate the overall performance loss of a partitioned cache.

The paper,

"Integrating Security Constraints into Fixed Priority Real-Time Schedulers"
by Sibin Mohan, Man-Ki Yoon, Rodolfo Pellizzoni, and Rakesh Bobba

touches upon an important new topic that has reached the design of real-time systems, vulnerability to security breaches. It assumes a system of tasks with mixed security levels where higher security tasks should be guarded against tasks with a lower security level which shall not be trusted. Whenever such tasks use the same physical resources, information might leak to the tasks with a lower security level if these tasks can evaluate the resource status. Several examples are given where such infringement only uses indirect information, such as cache misses. Therefore, "flushing" tasks are intermitted at any switch from higher to lower security level tasks. Such flushing entails overhead leading to a new scheduling problem. The journal paper motivates and formulates the security and scheduling requirements for the case of fixed priority scheduling and evaluates an optimization solution that improves over the ECRTS 2014 paper.

The paper

"Architecture and Analysis of a Dynamically-Scheduled Real-Time Memory Controller"
by Yonghui Li, Benny Akesson, and Kees Goossens

addresses a long-term challenge of real-time systems, the efficient use of DRAM memory with its highly stateful behavior. If the memory transactions have variable size, then optimizing memory timing is particularly difficult. The authors propose a memory controller that dynamically adapts memory command scheduling thereby handling systems with different classes of requests, guaranteed worst case response times for tasks with real-time requirements, and optimized average case response times for all other tasks. This is an important contribution to the field of mixed-criticality systems.

Please, enjoy the Special Issue!

Rolf Ernst is a professor at the Technische Universität Braunschweig, Germany, where he chairs the Institute of Computer and Network Engineering (IDA). His research covers embedded and cyberphysical systems with applications in automotive, avionics and smart buildings. Main aspects are real-time and safety critical systems. He is an IEEE Fellow and a member of the German Academy of Science and Engineering, acatech. He received the EDAA Lifetime Achievement Award 2014.