

## **Guest Editorial: From Uniprocessors to Multiprocessors: Advances in Real-Time Systems**

**Luis Almeida**

Published online: 18 May 2013  
© Springer Science+Business Media New York 2013

Real-time multiprocessor computing has been an active area of research for many years. For example, several real-time multiprocessor scheduling results are known for more than three decades. However, while the field of real-time uniprocessor computing evolved steadily to a mature area, the multiprocessor counterpart lagged behind, plagued by higher intrinsic complexity, costly platforms, low efficiency and few applications. This situation changed radically in the last decade with the massification of multi-core processors that made parallel computing platforms available in practically all general purpose computers as well as many industrial and embedded systems. This revolution forced adaptations in many existing software commodities to take advantage of the available parallel computing hardware and motivated a strong research effort in tackling the idiosyncrasies of these systems. The real-time systems community in particular has been dedicating substantial attention to issues such as the definition of analysable task models that better adapt to parallel platforms, the scheduling of the tasks on the set of processors, the tasks synchronization across processors and the access to global resources including caches, to name just a few.

This special issue includes four recent works that are representative of the state of the art in real-time multiprocessor computing. These are follow-ups of the four papers that were nominated for the Best and Student Best Paper Awards at RTSS 2011, the 32<sup>nd</sup> IEEE Real-Time Systems Symposium. Three of them address task scheduling issues while the fourth addresses worst-case execution time estimation. Moreover,

---

L. Almeida (✉)  
IT—FEUP, University of Porto, R. Dr. Roberto Frias, 4200-465 Porto, Portugal  
e-mail: [lda@fe.up.pt](mailto:lda@fe.up.pt)  
url: <http://www.fe.up.pt/~lda>

three of these works present analytical techniques that extend from uniprocessors to multiprocessors, contributing to a consolidation of the analysis of real-time systems.

In *Multi-core Real-Time Scheduling for Generalized Parallel Task Models*, Saifullah, Li, Agrawal, Lu and Gill, build upon a synchronous task model with tasks consisting of sequential segments each of which composed of an arbitrary number of parallel threads that execute to completion. They then propose a generalized task model in which parallel tasks can be decomposed into a set of sequential tasks and they show that such decomposition achieves a resource augmentation bound of 4 under global EDF (Earliest Deadline First) scheduling and 5 under partitioned DM (Deadline Monotonic) scheduling. In particular, the former bound is the first such bound derived for global EDF. They also show that such decomposition applies to parallel tasks generically modelled with DAG (Directed Acyclic Graphs) with unit execution requirement in each node.

In *Multiprocessor Scheduling by Reduction to Uniprocessor: an Original Optimal Approach*, Regnier, Lima, Massa, Levin and Brandt present the RUN (Reduction to UNiprocessor) optimal multiprocessor scheduler that reduces such scheduling problem to a series of uniprocessor problems. It is the first optimal multiprocessor scheduling algorithm that is not based on proportionate fairness, a technique that subdivides workloads to achieve similar deadlines for all tasks across the processors. Compared to proportionate fairness, RUN achieves optimality with significantly lower overhead for preemptions and migrations, reducing to partitioned EDF in all cases in which an adequate partitioning exists.

In *Resource Augmentation for Uniprocessor and Multiprocessor Partitioned Scheduling of Sporadic Real-Time Tasks*, Chen and Chakraborty provide new improved resource augmentation bounds using a linear approximation of the demand bound function (dbf). The improvements arise from considering the actual processor utilization violations with the approximate dbf. The new resource augmentation bounds are at most 1.6322 for uniprocessors and  $2.6322 - 1/M$  for constrained deadline tasks sets on uniform multiprocessors with  $M$  processors and  $3 - 1/M$  if the task set has arbitrary deadlines.

In *Scalable and Precise Refinement of Cache Timing Analysis via Path-sensitive Verification*, Chattopadhyay and Roychoudhury address the impact of the cache on the WCET (Worst-Case ExecutionTime) proposing a novel analytical framework that combines abstract interpretation and program verification, while considering different cache systems ranging from uniprocessors to multi-core processors. Different program verification techniques can be used, such as model checking and symbolic execution. The results achieved on top of the Chronos WCET analysis tool show significant improvements in precision with reasonable analysis execution time.

We believe these four papers are of great interest not only to specialists in each specific topic but also to unspecialized readers who wish a snapshot of the state of the art research in real-time multiprocessor systems.



**Luis Almeida** received a Licenciatura in Electronics and Telecommunications Engineering and a Ph.D. in Electrical Engineering, both from the University of Aveiro, Portugal, in 1988 and 1999, respectively. Since the end of 2008, he is an Associate Professor in the Electrical and Computer Engineering Department of the University of Porto, Portugal, and a senior researcher at IT and IEETA Research Units. His research interests include real-time networks for distributed industrial/embedded systems and control architectures for teams of mobile robots. He is a member of the IEEE Computer Society and the RoboCup Federation.