PREFACE

Special section on
Complex Problems in High-Performance Computing Systems

The possibilities offered by the ever increasing numbers of interconnected computing nodes in high performance systems continuously foster new applications that scale up in requirements and needs, together with the available computing power. Timely and accurate computing is required in most high value applications, but unfortunately scaling up does not in itself guarantee or favor such requirements, as expanding the number of nodes results in an increase in systems complexity, as well as in management and design problems. Issues that had a minor impact when the scale of commonly used parallel systems was still small, and manageable with heuristic approaches and manual settings, are now important challenges that must be faced with stable and formal theoretical methods, in order to guarantee what users need in system evolution.

This special section offers six research papers that present some interesting recent results on important aspects of the design and management of high-performance computing systems. The selection targets some issues regarding optimized management of computing resources, such as scheduling, communications, cooperation, synchronization and reliability in real applications or realistic model hypotheses.

Qureshi *et al.* deal with the problems in mapping the scheduling needs of a system with hardware scheduling capabilities, and specifically with the situation that arises when there is a mismatch between the cardinality of a periodic task set and hardware supported priority levels. The mapping affects the schedulability of the original task set, and the related properties have to be investigated as a result of the choice of given mapping criteria. The authors provide a mapping technique that keeps the properties of the original set and a feasibility test to support the decision process, together with formal verification and comparisons with previous works.

Bossard and Kaneko present a new topology to cope with the high numbers of nodes and connections in modern supercomputers, namely, torus-connected cycles. Their approach, based on a combination of a torus and a ring, allows better scalability while keeping a simple structure even with significantly high numbers of nodes, by maintaining all the advantages of the properties of torus networks and combining them with the benefits of having hierarchical interconnection networks. The authors present a formal analysis of their solution, showing the criteria to establish the diameter, a suitable point-to-point routing algorithm, and another one that solves the Hamiltonian path problem in two dimensions and the Hamiltonian cycle issue.

Gąsior and Seredyński provide an interesting solution for cloud computing environments that solves the problem of security-aware job scheduling. The authors propose an approach based on the Pareto dominance relationship on an individual user level and a subsequent equilibrium search process by modeling the provisioning of resources as a set of brokers representing the users and by targeting the allocation of a limited quantity of resources to a specific number of jobs, in order to minimize the probability of failing their execution and the total completion time to minimize both their execution failure probability and the total completion time. Their strategy exploits the *Spatial Prisoner's Dilemma* game-theoretic model to find the solution with a consistent decision process.

Mariano and Correia explore the problem of having cooperation in sharing an indivisible and durable resource by means of a game-theoretic approach. They analyze a multi-stage resource sharing game with two players to study in depth the dynamics of different coordination strategies based on implicit or explicit agreements. Their approach, which points out the existence of multiple Pareto optimal profiles in this game, constitutes a touchstone for investigating many real-world resource-sharing situations.

Feuerriegel and Bücker tackle a challenge and propose a technique to improve effective execution of the well-known Lanczos algorithm on systems composed of distributed memory computing nodes running a high number of processes. This algorithm, which is probably the principal tool used in a wide number of different application fields in engineering to compute the dominant eigenvalues of a large sparse non-symmetric matrix, suffers from a limitation due to the synchronization time, which significantly increases with the number of processes. The authors present some possible variants that reduce the need for synchronization, and compare their performances with previous approaches.

Zeifman *et al.* deal with the problem of service life duration assumptions in systems that, being affected by intense workloads or subject to stops for safety reasons, cannot be evaluated with the usual hypotheses of infinite life or stationarity. The authors offer a transient analysis method suitable to reliably model these situations, based on a class of Markov processes defined on non-negative integers, namely, inhomogeneous birth and death processes. The method is founded on the hypothesis that the system can jump back to the origin state from any of its states in a state-independent way, besides being able to regularly evolve between adjacent states with deterministic, eventually state-dependent, transition times, and is proved valid by means of numerical results.

We want to express our gratitude to the reviewers, who carefully examined the papers and eventually allowed us to produce, with their work and suggestions for the authors, a high-quality special session. In our opinion, the selected papers provide an interesting collection of results, methods and suggestions that can support the audience of the journal in their activities, and in general all scientists and professionals dealing with theoretical and practical research areas related to complex problems in distributed systems.

Mauro Iacono
Department of Political Science
Second University of Naples, Caserta, Italy

Joanna Kołodziej
Department of Computer Science
Cracow University of Tchnology, Poland

December 2015

Mauro Iacono received an M.Sc. degree in computer engineering from Università degli Studi di Napoli "Federico II" in 1999 and a Ph.D. degree in electrical engineering from Seconda Università degli Studi di Napoli in 2002. He works as a tenured assistant professor of computing systems at Seconda Università degli Studi di Napoli. He has published over 50 peer reviewed scientific papers and has served as a chairman, an IPC member and a referee for many conferences, as well as a guest editor, an editorial board member and a referee for several journals. His research is mainly focused on the field of performance modeling of complex computer-based systems, with special attention to multiformalism modeling techniques, and applications to critical systems and infrastructures, big data architectures and dependable systems. For more information, please visit `http://www.mauroiacono.com`.

Joanna Kołodziej graduated in mathematics from Jagiellonian University in Cracow in 1992, where she also received a Ph.D. in computer science in 2004. She works as an assistant professor at the Cracow University of Technology. She has been serving as a PC co-chair, a general co-chair and an IPC member of several international conferences and workshops, including *PPSN 2010, ECMS 2011, CISIS 2011, 3PGCIC 2011, CISSE 2006, CEC 2008, IACS 2008-2009, ICAART 2009-2010.* Dr. Kolodziej is an editorial board member and a guest editor of several peer-reviewed international journals, and an author and a co-author of many publications in high quality peer reviewed international journals. For more information, please visit `http://www.joannakolodziej.org/`.