J. Parallel Distrib. Comput. 100 (2017) 128–129

Contents lists available at ScienceDirect

J. Parallel Distrib. Comput.

journal homepage: www.elsevier.com/locate/jpdc

Editorial

Special issue on energy efficient multi-core and many-core systems, Part II

Multi-core and many-core architectures have become conventional to meet performance requirements of emerging applications ranging from massively parallel data centers to ultra-low power embedded devices for the Internet-of-things. It is expected that the number of cores in these systems increases dramatically in the near future. For such systems, energy efficiency is one of the primary design constraints. The cessation of Dennard scaling and the dark silicon phenomenon have limited recent improvements in transistor speed and energy efficiency, resulting in slowed improvements in multi-core and many-core systems. Consequently, architectural innovations have become crucial to achieve performance and efficiency gains. In addition, multi-core and many-core systems need to be able to reconfigure themselves adaptively by monitoring their own condition and the surrounding environment in order to adapt to different scenarios and performance-power requirements. Runtime resource management becomes crucial in modern parallel and distributed multicore systems due to increase in thermal issues as well as due to the need for various adaptive management techniques to maximize energy efficiency.

In this special issue, the Guest Editors have put together some of the new developments and trends in the context of energy efficient multi-core and many-core systems. We received fifty-three submissions and a total of six were accepted for this part. The high level of competition has led to the selection of top-level contributions covering a wide spectrum of topics in the domain of energy-efficient systems including on-chip communication, cache design, heterogeneity-supported model of computation, thermal management, and GPU-based acceleration.

As the number of cores on a single chip climbs towards hundreds, and even thousands, there is growing demand for energy-efficient and scalable on-chip interconnects. Nanophotonic interconnects enable new design possibilities for such many-core systems, however, they also require an off-chip laser source, which is often wasted to insertion losses and periods of low network activity. Kennedy et al., in their article entitled "CLAP-NET: Bandwidth Adaptive Optical Crossbar Architecture", present a nanophotonic crossbar architecture that uses decomposed optical crossbars and dynamic wavelength routing to manage bandwidth at runtime and maximize utility of the static laser power source.

In the same context, Tinati et al., in their article entitled "Topology Exploration of a Thermally Resilient Wavelength-Based ONoC", investigate how power consumed by processing elements located on electrical layer beneath the optical layer in 3D ICs affects thermally-induced network faults. They present an all-optical wavelength-based mesh architecture to address the effect of topology on operational faults.

Sardar et al., in their article entitled "Theorem Proving Based Formal Verification of Distributed Dynamic Thermal Management Schemes", propose a methodology based on theorem proving to perform formal verification of Distributed Dynamic Thermal Management (DDTM) techniques. The authors claim that using their methodology any number of cores and for all times can be specified and verified in terms of functional and timing properties. They present a generic higher-order-logic formalization model to formally specify any DDTM scheme.

Kee et al., in their article entitled "An Analytical Model Based on Performance Demand of Workload for Energy-Efficient Heterogeneous Multicore Systems", aim at extending the Amdahl's law to analyze the performance energy trade-off in the context of heterogeneous multicore systems. Heterogeneous multicore systems include cores with different performance power characteristics. For this reason, the authors propose an analytical model to study the benefits and proper configurations of the single instruction set architecture heterogeneous multicore system under realistic workload.

Wlotzka et al., in their article entitled "GPU-accelerated smoothers and grid transfer operators for parallel multigrid methods on HPC clusters", propose a parallel geometric multigrid linear solver using finite element discretizations of partial differential equations which runs on CPU- and GPU-based platform. Their aim is to experimentally address the question "whether the time to solution and energy to solution can benefit from using accelerators for smoothing and grid transfer." They present that, for a large problem size, using an accelerated asynchronous smoother can offer substantial savings of time and energy.

In the final article of this issue, an energy-efficient architectural design for cache memories is proposed by Valls *et al.* In their article entitled "The Tag Filter Architecture: An energy-efficient cache and directory design", they present the Tag Filter (TF) Architecture, a technique that can be applied to different types of set-associative caches. Their aim is to reduce the number of tags and data entries checked when accessing a target cache, in order to reduce the dynamic power consumption. TF Architecture enables the implementation of high-associativity caches with a similar energy consumption as that of conventional low-associativity caches, while reducing the number of conflict misses.

We sincerely hope the reader finds this special issue useful and that it will inspire further research in this very important area of

http://dx.doi.org/10.1016/j.jpdc.2016.10.009
0743-7315/

computer system design. We would like to thank all authors who submitted papers to this special issue. Special thanks go to the referees for their time and diligence during the review process and for providing us with high-quality reviews. Finally, we would like to thank Prof. Viktor Prasanna, Editor-in-Chief of the Elsevier Journal of Parallel and Distributed Computing, for offering us the opportunity to guest-edit this Special Issue.

Amir M. Rahmani
University of California, Irvine, USA
TU Wien, Austria

Pasi Liljeberg
University of Turku, Finland

Jose L. Ayala
Complutense University of Madrid, Spain

Hannu Tenhunen
KTH Royal Institute of Technology, Sweden

Alexander V. Veidenbaum
University of California, Irvine, USA

Available online 14 October 2016