Computing (2017) 99:725–726
DOI 10.1007/s00607-017-0570-9

EDITORIAL

Special issue on energy reduction techniques for exa-scale computing: theory and practice

Shajulin Benedict[1] · Michael Gerndt[2] · Siegfried Benkner[3]

Published online: 28 July 2017

Keywords AutoTuning · Energy Efficiency · Energy Prediction · Tools

1 Introduction

Recently, exa-scale computing has received much attention among the hardware and software designers although exa-scale machine will first become available in the 2020s. In fact, these machines are supposed to have hierarchical memory structures, highly scalable architectures with many accelerated nodes and a high performance IO system. Application and tool developers are rapidly developing solutions to address the known possible exa-scale challenges such as poor scalability, increased energy consumption, improper mapping of applications to hardware, and so forth. This special issue addresses the recent trends and technologies involved in reducing the energy consumption of scientific applications on future exa-scale systems.

✉ Shajulin Benedict
shajulin@sxcce.edu.in

Michael Gerndt
gerndt@in.tum.de

Siegfried Benkner
siegfried.benkner@univie.ac.at

1 HPCCLoud Research Laboratory, SXCCE, Anna University, Nagercoil 629003, India

2 Chair for Computer Architecture, Technische Universitaet Muenchen, Munich, Germany

3 Research Group Scientific Computing, University of Vienna, Vienna, Austria

2 The special issue: an overview

This special issue is based on the papers presented at the 2nd (IEEE) International Conference on Green High Performance Computing (ICGHPC'16) and the open call for papers announced at the journal site. We received 14 submissions towards this call; 8 papers were thoroughly peer-reviewed; and, 4 papers were selected. These four selected papers discuss static and dynamic tuning of the application's energy efficiency via DVFS, energy prediction for CUDA kernels, and techniques to more efficiently used GPGPUs as accelerator. Most experiments were conducted on supercomputers such as SuperMUC (Garching, Germany) and Tianhe-2 (Guangzhou, China). The glimpse of these papers is discussed below:

Paper 1: Schuchart et al present a tool based approach to improve the energy efficiency of scientific applications. It is obvious that the energy tuning of scientific applications will become a cumbersome task if manually handled. The authors have developed a tuning plugin based on the Periscope Tuning Framework as part of the READEX project—the European Union's Horizon H2020 project.

Paper 2: Guillen et al present a model-based tuning approach where the models are created from the performance data of HPC applications. The proposed autotuning approach was analyzed and evaluated at the Leibniz Supercomputing Centre, Germany.

Paper 3: A few dynamic regression model based prediction approaches were proposed by Rejitha et al. The approach targeted CUDA based applications on GPGPU architectures. The work was carried out at the HPCCLoud Research Laboratory, India, as part of the *Energy Tuning of Scientific Applications (EASE)* project—an Indo-Austrian DST-FWF funded project.

Paper 4: LU factorization kernel was designed with improved energy efficiency by Chen et al. In this paper, the authors consider a coprocessor resident implementation of LU in order to avoid hefty data transfers between CPU and GPUs. The work was experimentally evaluated at the Tianhe-2 supercomputer, China.

Acknowledgements The guest editors thank all the authors who have submitted their work to this Special Issue (SI); and, the reviewers of this SI for their timely reviews and constructive comments.