J Supercomput (2018) 74:527–529
https://doi.org/10.1007/s11227-018-2241-9

Special section: towards high performance computing in the cloud

George A. Gravvanis[1] · John P. Morrison[2] · Dan C. Marinescu[3] · Christos K. Filelis-Papadopoulos[1]

It gives us great pleasure to edit this special section on "Towards High Performance Computing in the Cloud," which contains papers from a number of high-profile authors in the field. These contributions were submitted via an open call for papers. All papers were rigorously reviewed and chosen on the basis of excellence and on the insight which they afford to the subject. As a result of the call for papers, twenty (20) articles were submitted and six (6) contributions were selected, resulting in an acceptance rate of 30%.

HPC workloads have unique requirements in terms of performance, resource requirements and storage. However, several issues, such as interconnection speeds, data locality, placement of VMs and partial support for specialized hardware, limit the effectiveness of HPC applications in the Cloud.

Nevertheless, in spite of ongoing challenges, recent years have seen HPC workloads migrating to the Cloud. This migration is driven by the desires to leverage novel

✉ George A. Gravvanis
ggravvan@ee.duth.gr

John P. Morrison
j.morrison@cs.ucc.ie

Dan C. Marinescu
dcm@cs.ucf.edu

Christos K. Filelis-Papadopoulos
cpapad@ee.duth.gr

1 Department of Electrical and Computer Engineering, School of Engineering, Democritus University of Thrace, University Campus, Kimmeria, 67100 Xanthi, Greece

2 Computer Science Department, University College of Cork, Cork, Ireland

3 Department of Electrical Engineering and Computer Science, University of Central Florida, Orlando, FL, USA

heterogeneous hardware, now being supported, and to flexibly access, the effectively unlimited, commodity resources available.

Extant service delivery models and resource allocation mechanisms associated with traditional Cloud environments are not designed for handling HPC applications. Therefore, tuning libraries to available heterogeneous hardware requires dedicated HPC experts and is difficult to perform in the generalized fashion required by the Cloud.

In their paper, *Large Scale Simulation of a Self-Organizing Self-Management Cloud Computing Framework*, Filelis-Papadopoulos et al. [1] discuss the large-scale simulation of a Self-Organizing and Self-Managing resource allocation system. The simulation of this heterogeneous system at scale required the development of various components in a parallel simulation framework. The authors consider the simulation of large-scale heterogeneous resources with two different VM strategies using proposed energy consumption and task execution models. Implementation details concerning the new functionalities of the parallel Cloud simulation framework are given, and simulation results obtained from executions of variable scale on a supercomputing environment, indicating performance and scalability, are given. Furthermore, discussions on the simulation of the Self-Organizing and Self-Managing system in relation to energy consumption and resource utilization are given.

In *Out-of-core Implementation for Accelerator Kernels on Heterogeneous Clouds*, by Khaleghzadeh et al. [2], a library (HCLOOC) is described which addresses issues, such as limited main memory of accelerators and limited bandwidth of the PCI-E bus, related to the execution of workloads in nodes containing accelerators such as GPUs, Xeon Phis or FPGAs. The framework is based on OpenCL command queues, Intel offload streams, and CUDA streams, allowing concurrent utilization of copy and execute engines. The authors elucidate the key features of their library using an out-of-core implementation of the matrix multiplication of large dense matrices on a hybrid node hosting all three of the aforementioned types of accelerators. The authors report that the given out-of-core implementation performs better than a state-of-the-art GPU implementation. Furthermore, no performance drop for the GPU, Xeon Phi and FPGA implementations was experienced when problem sizes exceeded the main memory size.

A distributed self-mesh adaptive N-body simulation algorithm is proposed by Kyziropoulos et al. The contribution is detailed in the article *Toward the design of a novel hybrid parallel N-body method in scope of modern cloud architectures* [3]. This hybrid parallel mesh-type method requires the solution of the Poisson equation in three space variables with boundary conditions obtained from multipole expansions. Shallow regions are adaptively refined to improve accuracy. The proposed scheme is parallelized using a uniform body distribution, requiring communications only during the computation of the potential and density distribution, and is based on approximate inverses and multiprojection techniques. The parallel performance of the scheme is evaluated in modern Cloud environments as well as supercomputing infrastructures, performing large-scale simulation of galaxies on high-resolution meshes. Moreover, the impact and potential of Cloud environments for N-body simulation in scope of performance is discussed.

The continuous growth of big data applications requires new ways of storing heterogeneous datasets. Sandhu et al., in their contribution, *TDRM: Tensor-based Data*

Representation and Mining for Healthcare Data in Cloud Computing Environments [4], discuss the topic of storing heterogeneous data in tensor format. The tensors store data in a raw format, and these are replicated to other tensors at different abstraction layers. Moreover, details of a mathematical foundation of tensor formation and query processing are given. The proposed system performs better than existing approaches, in terms of response time, accuracy and relevancy.

The concept of elastic online scheduling of big data streaming applications with respect to high-velocity continuous data streams is discussed by Sun et al. on their manuscript *Rethinking Elastic Online Scheduling of Big Data Streaming Applications over High Velocity Continuous Data Streams* [5]. The authors proposed E-stream, an elastic online scheduling framework for big data streaming applications based on profiling mathematical relationships between system response time, multiple application fairness and online features of high velocity. The framework scales in and out of data stream graphs quantifying computation and communication and scheduling the graphs based on max–min fairness strategy. The proposed system is evaluated using real-world big data system and is compared to the Storm framework.

Sung et al. in their article *OLM: Online LLC Management for Container-based Cloud Service* discuss the problem of last-level cache (LLC) sharing in container-based Cloud environments [6]. The authors propose a scheme for LLC sharing based on indirect LLC usage pattern obtained from monitoring the memory access pattern. The scheme separates executing applications into two groups using offline profiling data at runtime: (a) cache-thrashing applications, which fill up LLC with data with no temporal locality, and (b) non-cache-thrashing applications. These groups use different partitions of the LLC and are isolated using OS-based page coloring. Moreover, the proposed scheme improves container performance by up to 40%.

We hope that the reader will find the papers in this issue to be both informative and inspiring. We are most grateful to and would like to thank all of the authors for sharing their insights and experiences. Additionally, special thanks are due to the reviewers for their time, effort and valuable criticism and suggestions.

References

1. Filelis-Papadopoulos CK, Giannoutakis KM, Gravvanis GA, Tzovaras D (2017) Large scale simulation of a self-organizing self-management cloud computing framework. J Supercomput. https://doi.org/10.1007/s11227-017-2143-2
2. Khaleghzadeh H, Zhong Z, Reddy R, Lastovetsky A (2017) Out-of-core implementation for accelerator kernels on heterogeneous clouds. J Supercomput. https://doi.org/10.1007/s11227-017-2141-4
3. Kyziropoulos PE, Filelis-Papadopoulos CK, Gravvanis GA, Efthymiopoulos C (2017) Towards the design of a novel hybrid parallel N-body method in scope of modern cloud architectures. J Supercomput. https://doi.org/10.1007/s11227-017-2140-5
4. Sandhu R, Kaur N, Sood SK, Buyya R (2017) TDRM: tensor-based data representation and mining for healthcare data in cloud computing environments. J Supercomput. https://doi.org/10.1007/s11227-017-2163-y
5. Sun D, Yan H, Gao S, Liu X, Buyya R (2017) Rethinking elastic online scheduling of big data streaming applications over high velocity continuous data streams. J Supercomput. https://doi.org/10.1007/s11227-017-2151-2
6. Sung H, Kim M, Min J, Eom H (2017) OLM: online LLC management for container-based cloud service. J Supercomput. https://doi.org/10.1007/s11227-017-2181-9