

The 5th International Workshop on Accelerators and Hybrid Exascale Systems (AsHES'15)

As we look beyond the petascale era, accelerator and heterogeneous architectures are expected to play a preeminent role in architecting the largest systems in the world. Future systems may not only offer accelerators (e.g. GPGPU, Intel® Xeon Phi™, FPGA) and hybrid processors of both lightweight and heavyweight cores (e.g. APU, big.LITTLE), but also may use hybrid memory systems equipped with stacked memory and non-volatile memory in addition to regular DRAM. While there is significant interest in these architectures, much of it is an artifact of the hype associated with them. This workshop focuses on understanding the implications of accelerators and heterogeneous designs on the hardware systems, applications, and programming environments of future systems. It seeks to ground accelerator research through studies of application kernels or whole applications on such systems, as well as tools and libraries that improve the performance and productivity of applications on these systems.

The goal of this workshop is to bring together researchers and practitioners who are involved in application studies for accelerators and heterogeneous systems, to learn the opportunities and challenges in future design trends for HPC applications and systems. The Program Committee of 18 members completed a total of 42 reviews, followed by online discussions. The final program selected 7 papers, by authors from 6 countries, presenting frontier research in hybrid systems design, optimization, and application.

Workshop Program:

8:30-8:45 Opening remarks

8:45-9:50 Keynote by Michela Taufer -- The Numerical Reproducibility Fair Trade: Facing the Concurrency Challenges at the Extreme Scale

9:50-10:30 Break

10:30-12:00 Session 1: Accelerating Analytics

1. Sina Meraji, Sunil Kamath, John Keenleyside and Bob Blainey, Towards A Combined Grouping and Aggregation Algorithm for Fast Query Processing in Columnar Databases with GPUs
2. Dipanjan Sengupta, Kapil Agarwal, Shuaiwen Song and Karsten Schwan, GraphReduce: Large-Scale Graph Analytics on Accelerator-Based HPC Systems
3. Shuai Che, Gregory Rodgers, Brad Beckmann and Steve Reinhardt, Graph Coloring on the GPU and Some Techniques to Improve Load Imbalance

12:00-1:30 Lunch

1:30-3:30 Session 2: Algorithm Design for Heterogeneous Systems

1. Sushil K. Prasad, Michael McDermott, Xi He and Satish Puri, GPGPU-based Parallel R-tree Construction and Querying
2. Aditya Deshpande and P J Narayanan, Fast Burrows Wheeler Compression Using All-Cores
3. Kiran Raj Ramamoorthy, Dip Sankar Banerjee, Kannan Srinathan and Kishore Kothapalli, A Novel Heterogeneous Algorithm for Multiplying Scale-Free Sparse Matrices
4. Kazuya Matsumoto, Toshihiro Hanawa, Yuetsu Kodama, Hisafumi Fujii and Taisuke Boku, Implementation of CG Method on GPU Cluster with Proprietary Interconnect TCA for GPU Direct Communication

Organizers:

General Chair

Xiaosong Ma, Qatar Computing Research Institute, Qatar; North Carolina State University, USA

Program Chairs

James Dinan, Intel Corporation, USA

Wenguang Chen, Tsinghua University, China

Steering Committee

Pavan Balaji, Argonne National Laboratory, USA (co-chair)

Jiayuan Meng, Argonne National Laboratory, USA (co-chair)

Satoshi Matsuoka, Tokyo Institute of Technology, Japan

Yunquan Zhang, Chinese Academy of Sciences, China

Xiaosong Ma, Qatar Computing Research Institute, Qatar; North Carolina State University, USA

Publicity Chairs

Jeff R. Hammond, Intel Labs, USA

Technical Program Committee

Surendra Byna, Lawrence Berkeley National Laboratory, USA

Huimin Cui, Chinese Academy of Sciences, China

Rong Ge, Marquette University, USA

Junli Gu, AMD, China

Khaled Hamidouche, The Ohio State University, USA

Siva Hari, NVidia, USA

Feng Ji, Google, USA

Gabriele Jose, Intel Corporation, USA

Sriram Krishnamoorthy, Pacific Northwest National Laboratory, USA

Naoya Maruyama, RIKEN Advanced Institute for Computational Science, Japan

Antonio J. Peña, Argonne National Laboratory, USA

Kelly Shaw, University of Richmond, USA

Xipeng Shen, College of William and Mary, USA

Guanyu Sun, Peking University, China

Bronis de Supinski, Lawrence Livermore National Laboratory, USA

Hao Wang, Virginia Tech, USA

Fangfang Xia, Argonne National Laboratory, USA

Yongpeng Zhang, Stone Ridge Technology, USA

Additional Reviewers

Bin Ren, Pacific Northwest National Laboratory, USA

Bo Wu, Colorado School of Mines, USA

Qi Zhu, National University of Defense Technology, China

Acknowledgment

We thank the authors for submitting their work to our workshop. We are also grateful to our PC members who help us maintain the high quality of the workshop. We further acknowledge the input from external reviewers for their effort.