The 3rd International Workshop on Accelerators and Hybrid Exascale Systems (AsHES'13)

As we look beyond the petascale era, accelerators such as Graphics Processing Units (GPUs) and FPGAs, as well as upcoming integrated hybrid processing cores, are expected to play a preeminent role in architecting the largest systems in the world. While there is significant interest in these architectures, much of this interest is an artifact of the hype associated with them. This workshop focuses on understanding the implications of accelerators on the architectures 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 or productivity of applications trying to use these systems.

The goal of this workshop is to bring together researchers and practitioners who are involved in application studies for accelerators and other hybrid systems, to learn the opportunities and challenges in future design trends for HPC applications and systems. This year, AsHES has continued its growth since it was established in 2011, and received a record 24 submissions. The Program Committee of 18 members completed a total of 105 reviews, followed by online discussions. The final program selected 12 papers, by authors from 6 countries, presenting frontier research in hybrid systems design, optimization, and application.

Workshop Program:

8:30-8:45 Opening remarks
8:45-9:45 Keynote by Richard Vuduc
9:45-10:15 Break
10:15-11:55 Session 1: Programing Model and Performance Optimizations

1. Ashwin Aji, Pavan Balaji, James Dinan, Wuchun Feng and Rajeev Thakur. Synchronization and Ordering Semantics in Hybrid MPI+GPU Programming
2. Toshihiro Hanawa, Yuetsu Kodama, Taisuke Boku and Mitsuhisa Sato. Tightly Coupled Accelerators Architecture for Minimizing Communication Latency among Accelerators
3. Yash Ukidave and David Kaeli. Analyzing Optimization Techniques for Power Efficiency on Heterogeneous Platforms
4. Andra-Ecaterina Hugo, Abdou Guermouche, Pierre-Andre Wacrenier and Raymond Namyst. Composing multiple Starpu applications over heterogeneous machines: a supervised approach

11:55-1:30 Lunch
1:30-3:10 Session 2: Accelerated Applications

1. Dip Sankar Banerjee, Parikshit Sakurikar and Kishore Kothapalli. Fast, Scalable Parallel Comparison Sort on Hybrid Multicore Architectures
2. Ichitaro Yamazaki. Tridiagonalization of a symmetric dense matrix on a GPU cluster
3. Guanghao Jin, Toshio Endo and Satoshi Matsuoka. A Multi-level Optimization Method for Stencil Computation on the Domain that is Bigger than Memory Capacity of GPU
4. Yang You, Haohuan Fu, Guangwen Yang and Xiaomeng Huang. Accelerating the 3D Elastic Wave Forward Model on GPU and MIC

3:10-3:40 Break

978-0-7695-4979-8/13 $26.00 © 2013 IEEE
DOI 10.1109/IPDPSW.2013.290

3:40-5:20 Session 3: Emerging Hybrid Systems

1. Michael Boyer, Jiayuan Meng and Kalyan Kumaran. Improving GPU Performance Prediction with Data Transfer Modeling
2. Gaurav Mitra, Beau Johnston, Alistair P. Rendell, Eric McCreath and Jun Zhou. Use of SIMD Vector Operations to Accelerate Application Code Performance on Low-Powered ARM and Intel Platforms
3. Gao Tao, Lu Yutong and Suo Guang. Using MIC to accelerate a typical data-intensive application: the Breadth-first Search
4. Robert Overman, Jan Prins, Laura Miller and Michael Minion. Dynamic Load Balancing of the Adaptive Fast Multipole Method in Heterogeneous Systems

Organizers:

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

General Chair

Jiayuan Meng, Argonne National Laboratory, USA

Program Chairs

Yunquan Zhang, Chinese Academy of Sciences, China
Xiaosong Ma, North Carolina State University, USA

Publicity Chairs

Weisong Shi, Wayne State University, USA

Web Chair
Feng Ji, North Carolina State University, USA
Liang Yuan, Chinese Academy of Sciences, China

Technical Program Committee

Hong An, University of Science and Technology of China, China
David Bader, Georgia Institute of Technology, USA
Surendra Byna, Lawrence Berkeley National Laboratory, USA
Ali R. Butt, Virginia Tech, USA
Wenguang Chen, Tsinghua University, China
Xiaobing Feng, Institute of Computing Technology, CAS, China
Torsten Hoefler, ETH Zürich., Switzerland
Heshan Lin, Virginia Tech, USA
Yutong Lu, National University of Defense Technologies, China
Naoya Maruyama, Tokyo Institute of Technology, Japan
Frank Mueller, North Carolina State University, USA
Dimitris Nikolopoulos, University of Crete/FORTH, Greece
P. Sadayappan, Ohio State University, USA
Xipeng Shen, College of William and Mary, USA
Bronis de Supinski, Lawrence Livermore National Laboratory, USA
Vinod Tipparaju, Advanced Micro Devices, Inc., USA

Dawei Wang, Illinois Institute of Technology, USA
Fangfang Xia, Argonne National Laboratory, USA

Journal Special Issue Co-editors

Jiayuan Meng, Argonne National Laboratory
Toshio ENDO, Tokyo Institute of Technology

Additional Reviewers
Ashwin Aji, Virginia Tech, USA
Amir Bahmani, North Carolina State University, USA
Filip Blagojevic, Hitachi Global Storage Technologies, USA
Shuai Che, AMD Research, USA
Li Chen, Institute of Computing Technology, CAS, China
Benjamin Clay, North Carolina State University, USA
Cong Du, Arista Networks, Inc., USA
James Elliott, North Carolina State University, USA
David Fiala, North Carolina State University, USA
Oded Green, Georgia Institute of Technology, USA
Feng Ji, North Carolina State University, USA
Konstantinos Krommydas, Virginia Tech, USA
Dong Li, Oak Ridge National Laboratory, USA
Min Li, Virginia Tech, USA
Ying Liu, Institute of Computing Technology, CAS, China
Fang Lu, Institute of Computing Technology, CAS, China
Spyros Lyberis, Institute of Computer Science, Foundation for Research and Technology - Hellas, Greece
Fei Meng, North Carolina State University, USA
Carlo Mundo, Virginia Tech, USA
Vassilis Papaefstathiou
Antonio Peña, Universitat Jaume I, Spain
Miquel Pericàs, Tokyo Institute of Technology, Japan
M. Mustafa Rafique, Virginia Tech, USA
Vignesh Ravi, AMD, USA
Arash Rezaei, North Carolina State University, USA
Tom Scogland, Virginia Tech, USA
Zhiming Shen, North Carolina State University, USA
Kai Tian, Microsoft, USA
Thomas Uram, Argonne National Laboratory, USA
Mohamed Wahib, RIKEN Advanced Institute for Computational Science, Japan
Zhaoming Yin, Georgia Institute of Technology, USA
Jing Zhang, Virginia Tech, USA
Yao Zhang, University of Chicago, USA
Yongpeng Zhang, Stone Ridge Technology, USA
Zhijia Zhao, College of William and Mary, USA

Acknowledgement:

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.