

GPGPU-6 Final Program

March 16, 2013

Houston, TX

8:30 Welcome:

8:45-9:30 Keynote: Intel® Xeon Phi™ programmability: the good, the bad and the ugly. Robert Geva, Principle Engineer, parallel language architect, Intel.

9:30-10:30 Optimizations

- Reducing Divergence in GPGPU Programs with Loop Merging, Tianyi David Han and Tarek Abdelrahman.
- Split Tiling for GPUs: Automatic Parallelization Using Trapezoidal Tiles to Reconcile Parallelism and Locality, avoiding Divergence and Load Imbalance, Albert Cohen, Tobias Grosser, Paul Kelly, J Ramanujam, P Sadayappan and Sven Verdoolaege.
- Memory Reuse Optimizations in the R-Stream Compiler, Nicolas Vasilache, Muthu Baskaran, Benoit Meister and Richard Lethin.

10:30-11:00 Break

11:00-12:20 Tools and Frameworks

- Valar: A Benchmark Suite to Study the Dynamic Behavior of Heterogeneous Systems, Perhaad Mistry, Yash Ukidave, Dana Schaa and David Kaeli.
- OpenCL C++, Benedict Gaster and Lee Howes.
- Fast Dynamic Memory Allocator for Massively Parallel Architectures, Sven Widmer, Dominik Wodniok, Nicolas Weber and Michael Goesele.
- Formalizing Address Spaces with application to Cuda, OpenCL, and beyond, Benedict Gaster and Lee Howes.

12:20-2:00 Lunch

2:00-2:40 Autotuning

- Warp Size Impact in GPUs: Large or Small? Ahmad Lashgar, Amirali Baniyasi and Ahmad Khonsari.
- Input-Aware Auto-Tuning for Directive-based GPU Programming, Alberto Magni, Dominik Grewe and Nick Johnson.

2:40-3:00 Break

3:00-4:00 Algorithms

- Atomic-free Irregular Computations on GPUs, Rupesh Nasre., Martin Burtscher and Keshav Pingali.
- Comparison Based Sorting for Systems with Multiple GPUs, Ivan Tanasic, Lluís Vilanova, Marc Jorda, Javier Cabezas, Isaac Gelado, Nacho Navarro and Wen-mei Hwu.
- Betweenness Centrality on GPUs and Heterogeneous Architectures, Ahmet Erdem Sariyüce, Kamer Kaya, Erik Saule and Umit V. Catalyurek.

4:00-5:00 Applications

- Accelerating Simulation of Agent-Based Models on Heterogeneous Architectures, Jin Wang, Norman Rubin, Haicheng Wu and Sudhakar Yalamanchili.
- Accelerating Financial Applications on the GPU, Scott Grauer-Gray, William Killian, Robert Searles and John Cavazos.
- Exploring GPU Architectures To Accelerate Semantic Comparison For Intention-Based Search, Ozgur Gonen, Jaskirat Batra, Sonali Mahapatra and Steve Liu.