

GPGPU-8

February 7, 2015

San Francisco, CA

8:30 – Welcome:

8:40-9:30 – Keynote: *“Dark secrets of heterogeneous memory models”*
Lee Howes, Qualcomm

9:30-10:00 - HPC

- Ayman Tarakji, Lukas Börger and Rainer Leupers. A Comparative Investigation of Device-Specific Mechanisms for Exploiting HPC Accelerators

10:00-10:30 - break

10:30-12:00 – Caching and Shared Memory

- Javier Cabezas, Marc Jordà, Isaac Gelado, Nacho Navarro and Wen-Mei Hwu. GPU-SM: shared memory multi-GPU programming
- Yingying Tian, Sooraj Puthoor, Joseph L. Greathouse, Bradford M. Beckmann and Daniel A. Jiménez. Adaptive GPU Cache Bypassing
- Mahmoud Khairy, Amr G. Wassal and Mohamed Zahran. Efficient Utilization of GPGPU Cache Hierarchy

12:00-14:00 – Lunch

14:00-15:30 – Optimization

- Jared Coplin and Martin Burtscher. Effects of Source-Code Optimizations on GPU Performance and Energy Consumption
- Azzam Haidar, Tingxing Dong, Piotr Luszczek, Stanimire Tomov and Jack Dongarra. Optimization for Performance and Energy for Batched Matrix Computations on GPUs
- Thibaut Lutz, Christian Fensch and Murray Cole. Helium: a Transparent Inter-kernel Optimizer for OpenCL

15:30-16:00 – break

16:00-18:00 – Applications

- Rashid Kaleem, Sreepathi Pai and Keshav Pingali. Stochastic Gradient Descent on GPUs
- Leiming Yu, Yan Zhang, Xiang Gong, Nilay Roy, Lee Makowski and David Kaeli. High Performance Computing of Fiber Scattering Simulation
- Molly A. O’Neil and Martin Burtscher. Rethinking the Parallelization of Random-Restart Hill Climbing: A Case Study in Optimizing a 2-Opt TSP Solver for GPU Execution
- Mahesh Ravishankar, Justin Holewinski and Vinod Grover. Forma: A DSL for Image Processing Applications to target GPUs and multi-core CPUs