GPGPU-10
Austin, TX

February 5, 2017

Final Program

7:30-8:30 – Breakfast Available

8:30 – Welcome: The Organizers

8:30-9:30 – Keynote 1 - Greg Stoner, Sr. Director in Radeon Open Compute, Chairman of the Board of the HSA Foundation

- *"ROCm: crafting an open platform for GPU computing exploration and problem solving,"*

9:30-10:00 - Embedded GPUs

- *"Efficient Convex Optimization on GPUs for Embedded Model Predictive Control,"* Leiming Yu, Abraham Goldsmith and Stefano DiCairano, Mitsubishi Electric Research Labs (MERL) and Northeastern University

10:00-10:30 – Coffee Break

10:30-12:00 – GPU Benchmarking and Optimization

- "*"High-Performance Cholesky Factorization for GPU-only Execution,"* Azzam Haidar, Ahmad Abdelfattah, Stanimire Tomov and Jack Dongarra, University of Tennessee
- *"Parallel CCD++ on GPUs for Matrix Factorization,"* Israt Nisa, Aravind Sukumaran-Rajam, Rakshith Kunchum and P. Sadayappan, Ohio St. University
- *"DNNMark: A Deep Neural Network Benchmark Suite for GPUs,"* Shi Dong and David Kaeli, Northeastern University

12:00-13:30 – Lunch (on your own)

13:30-14:30 – Keynote II - Hars Vardhan, Staff Research Engineer, Computer Science Innovation Center, Samsung Research America

- *"Effectively Scaling out Deep Learning Frameworks with GPUs,"*

14:30-15:00 – GPU Security

- *"Understanding the Security of Discrete GPUs,"* Zhiting Zhu, Sangman Kim, Yuri Rozhanski, Yige Hu, Emmett Witchel and Mark Silberstein, University of Texas at Austin and the Technion

15:00-15:30 – Coffee Break

15:30-17:00 – Compiler Optimizations

- *"Directive-based Tile Abstraction to Distribute Loops on Accelerators,"* Tristan Vanderbruggen, John Cavazos, Chunhua Liao and Daniel Quinlan, University of Delaware and Lawrence Livermore National Labs
- *"Launch-Time Optimization of OpenCL GPU Kernels,"* Andrew Lee and Tarek Abdelrahman, University of Toronto
- "*Multi-job Scheduling for OpenCL Kernels on CPU/GPU Platforms,"* Yuan Wen and Michael O'Boyle, University of Edinburgh