

GPGPU-11

Vienna, Austria

February 25, 2018
Final Program

8:00-8:30 – Breakfast Available

8:30 – Welcome: The Organizers

8:30-9:30 – Keynote 1 – Jun Yang, William Kepler Whiteford Professor of Electrical and Computer Engineering, University of Pittsburgh

- Initial Steps toward Making GPU a First-Class Computing Resource: Sharing and Resource Management

9:30-10:00 - Persistent Data Structures

- **A Case For Persist Barriers in GPUs**, Dibakar Gope, Arkaprava Basu, Sooraj Puthoor and Mitesh Meswani, ARM Research and AMD Research

10:00-10:30 – Coffee Break

10:30-12:00 – Applications/Frameworks

- **Overcoming the Difficulty of Large-scale CGH Generation on multi-GPU Cluster**, Takanobu Baba, Shinpei Watanabe, Boaz Jessie Jackin, Takeshi Ohkawa, Kanemitsu Ootsu, Takashi Yokota, Yoshio Hayasaki and Toyohiko Yatagai, Utsunomiya University and National Institute of Information and Communications Technology
- **Transparent Avoidance of Redundant Data Transfer on GPU-enabled Apache Spark**, Ryo Asai, Masao Okita, Fumihiko Ino and Kenichi Hagihara, Osaka University
- **GPU-based Acceleration of Detailed Tissue-Scale Cardiac Simulations**, Neringa Altanaite and Johannes Langguth, Simula Research Laboratory, Norway

12:00-13:30 – Lunch (on your own)

13:30-14:30 – Keynote 2 – Christophe Dubach, University of Edinburgh

- Generating High Performance GPU Code using Rewrite Rules with Lift

14:30-15:00 – Coffee Break

15:30-16:30 – Concurrent Kernels

- **MaxPair: Enhance OpenCL Concurrent Kernel Execution by Weighted Maximum Matching**, Yuan Wen, Michael O'Boyle and Christian Fensch, Trinity College Dublin, University of Edinburgh, Heriot-Watt University
- **Oversubscribed Command Queues in GPUs**, Sooraj Puthoor, Xulong Tang, Joseph Gross and Bradford Beckmann, AMD Research, Penn St. University and ARM