

Workshop on General Purpose Processing on Graphics Processing Units

October 4, 2007

FINAL PROGRAM

8:00 **Welcome and keynote introduction** – David Luzzi (NU)

8:10-9:00 **Keynote:** “GP Computing: Hardware, Architecture, Tools and Education,” Wen-mei Hwu, University of Illinois Urbana-Champaign.

9:00-10:20 **Session 1: Parallelizing Matrix/Lattice Operations** - Session Chair: Miriam Leeser (NU)

- “General-Purpose Sparse Matrix Building Blocks using the NVIDIA CUDA Technology Platform,” M. Christen, O. Schenk and H. Burkhardt, University of Basel, Switzerland.
- “Exploiting Structure of Symmetric or Triangular Matrices on a GPU,” J.-H. Jung and D. O’leary, University of Maryland, College Park.
- “How GPUs Can Improve the Quality of Magnetic Resonance Imaging,” S. Stone, H. Yi, J. Haldar, W. Hwu, B. Sutton and Z.-P. Liang, University of Illinois at Urbana-Champaign.
- “Fine-Grained Parallelization of Lattice QCD Kernel Routine on GPU,” K. Ibrahim and F. Bodin, IRISA/INRIA, Rennes, France, O. Pene, LPT/CNRS, Orsay, France.

10:20-10:40 **Break**

10:40-12:00 **Session 2: Power and Performance** - Session chair: David Luebke (NVIDIA)

- “PowerRed: A Flexible Power Modeling Framework for Power Efficiency Exploration in GPUs,” K. Ramani, University of Utah, A. Ibrahim, and D. Shimizu, AMD.
- “ENCORE: Energy-Conscious Rendering for Mobile Devices,” C.-H. Chang, P. Lohrmann, E. Agu and R. Lindeman, Worcester Poly Tech.
- “A Performance Study of General Purpose Applications on Graphics Processors,” S. Che, J. Meng, J. Sheaffer, and K. Skadron, University of Virginia.
- “Program Optimization Study on a 128-Core GPU,” S. Ryoo, C. Rodrigues, S. Stone, S. Baghsorkhi, S.-Z. Ueng and W. Hwu, University of Illinois Urbana-Champaign.

12:00-1:00 **Lunch and View Demos**

1:00-3:00 **Industry Session** - Session Chair: David Kaeli (NU)

- “Advances in Real-Time Medical Visualization,” Natalya Tatarchuk, AMD.
- “System Optimization of the Cell B.E. Heterogeneous Chip Multiprocessor for Efficient and Flexible Programmability,” Michael Gschwind, IBM Research.
- “NVIDIA’s CUDA: Democratizing Parallel Computing,” David Luebke, NVIDIA.

3:00-3:20 **Break**

3:20-4:20 **Session 3: Fast Algorithms** - Session Chair: Norm Rubin (ATI)

- “Fast Parallel GPU-Sorting Using a Hybrid Algorithms,” E. Sintorn and U. Assarsson, Chalmers University, Sweden.
- “On Implementing Graph Cuts on CUDA,” M. Hussein, A. Varshney and L. Davis, University of Maryland.
- “Stream-Based LDPC Decoding on GPUs,” G. Falcao and V. Silva, University of Coimbra, Portugal, S. Yamagiwa and L. Sousa, University of Lisbon, Portugal.

4:20-4:35 **Break**

4:35-5:15 **Session 4: Portability and Streaming** - Session Chair: Lalit Jain (Angstrom)

- “HMPP: A Hybrid Multi-core Parallel Programming Environment,” R. Dolbeau, S. Bihan and F. Bodin, CAPS Enterprise, France.
- “Efficient Stream Reduction on the GPU,” D. Rogers, Grenoble University, France, U. Assarsson, Chalmers University, Sweden, N. Holzschuch, Cornell.

5:15 **Workshop Concludes**