IA^3 2014
Workshop on Irregular Applications: Architectures and Algorithms

In cooperation with SIGHPC
Held in conjunction with SC14

Sunday, November 16
New Orleans Convention Center
Room: 273

8:50 – 9:00 Welcome – Antonino Tumeo (PNNL), John Feo (PNNL), Oreste Villa (NVIDIA)

9:00 – 10:00 Keynote 1: Prof. *Onur Mutlu,* Carnegie Mellon University
Chair: Antonino Tumeo (PNNL)

Rethinking Memory System Design (for Data-Intensive Computing)

Abstract:
The memory system is a fundamental performance and energy bottleneck in almost all computing systems. Recent system design, application, and technology trends that require more capacity, bandwidth, efficiency, and predictability out of the memory system make it an even more important system bottleneck. At the same time, DRAM and flash technologies are experiencing difficult technology scaling challenges that make the maintenance and enhancement of their capacity, energy-efficiency, and reliability significantly more costly with conventional techniques.

In this talk, we examine some promising research and design directions to overcome challenges posed by memory scaling. Specifically, we discuss three key solution directions: 1) enabling new memory architectures, functions, interfaces, and better integration of the memory and the rest of the system, 2) designing a memory system that intelligently employs multiple memory technologies and coordinates memory and storage management using non-volatile memory technologies, 3) providing predictable performance and QoS to applications sharing the memory/storage system. If time permits, we may also briefly describe our ongoing related work in combating scaling challenges of NAND flash memory.

An accompanying short paper, slightly outdated, can be found here:
http://users.ece.cmu.edu/~omutlu/pub/memory-scaling_memcon13.pdf

Bio:
Onur Mutlu is the Strecker Early Career Professor at Carnegie Mellon University. His broader research interests are in computer architecture and systems, especially in the interactions between languages, operating systems, compilers, and microarchitecture. He enjoys teaching and researching problems in computer architecture, including those related to the design of memory/storage systems, multi-core architectures, and scalable and efficient systems. He obtained his PhD and MS in ECE from the University of Texas at Austin (2006) and BS degrees in Computer Engineering and Psychology from the University of Michigan, Ann Arbor. Prior to Carnegie Mellon, he worked at Microsoft Research (2006-2009), Intel Corporation, and Advanced Micro Devices. He was a recipient of the IEEE Computer Society Young Computer Architect Award, Intel Early Career Faculty Honor Award, Faculty partnership Awards from IBM, HP, and Microsoft, a number of best paper awards, and a number of "computer architecture top pick" paper selections by the IEEE Micro magazine. For more information, please see his webpage at http://www.ece.cmu.edu/~omutlu.

10:00 – 10:30 Break

10:30 – 11:20 Keynote 2: Prof. *Keshav Pingali,* the University of Texas at Austin
Chair: David Haglin (PNNL)
Graph Analytics on the Galois System

Abstract:
Unstructured data is being generated at a tremendous rate in modern applications as diverse as social networks, recommender systems, genomics, health care and energy management. Networks are an important example of unstructured data and may arise explicitly, as in social networks, or implicitly, as in recommender systems.
These networks are challenging to handle; not only are they large-scale but they are constantly evolving, and many applications require difficult prediction tasks to be solved, such as link or ratings prediction.

In this talk, we will describe the Galois system, which is a data-centric parallel programming system that provides an easy-to-use programming model for implementing and tuning parallel irregular applications such as graph analytics algorithms. We will present performance numbers for implementations of some of these algorithms on the Galois system from our group and from other groups in academia and industry.

Bio:
Keshav Pingali is a Professor in the Department of Computer Science at the University of Texas at Austin, and he holds the W.A."Tex" Moncrief Chair of Computing in the Institute for Computational Engineering and Sciences (ICES) at UT Austin. He was on the faculty of the Department of Computer Science at

Cornell University from 1986 to 2006, where he held the India Chair of Computer Science.
Pingali's research has focused on programming languages and compiler technology for program understanding, restructuring, and optimization.
His group is known for its contributions to memory-hierarchy optimization; some of these have been patented and are in use in industry compilers.

Pingali is a Fellow of the ACM, IEEE and AAAS. He was the co-Editor-in-chief of the ACM Transactions on Programming Languages and Systems, and currently serves on the editorial boards of the International Journal of Parallel Programming and Distributed Computing. He also served on the NSF CISE Advisory Committee (2009-2012).

11:20 – 12:05 Paper Session 1: Applications and algorithms (including short papers) – Chair: Vito Giovanni Castellana (PNNL)

Using Vertex-Centric Programming Platforms to Implement SPARQL Queries on Large Graphs (25 min)
Eric L. Goodman (Sandia National Laboratories), Dirk Grunwald (University of Colorado)

An Extremely Fast Algorithm for Identifying High Closeness Centrality Vertices in Large-Scale Networks (10 min)
Vladimir Ufimtsev (University of Nebraska at Omaha) , Sanjukta Bhowmick (University of Nebraska at Omaha)

Heterogeneous Concurrent Execution of Monte Carlo Photon Transport on CPU, GPU and MIC. (10 min)
Noah Wolfe (Rensselaer Polytechnic Institute), Tianyu Liu (Rensselaer Polytechnic Institute), Christopher Carothers (Rensselaer Polytechnic Institute), and Xie George Xu (Rensselaer Polytechnic Institute)

1:30 – 3:00 Panel Session: Architectures and Algorithms for Irregular Applications
Moderator: John Feo (PNNL)

Panelists: *Clayton Chandler* (DOD), *Benoit Dupont de Dinechin* (Kalray), *Maya Gokhale* (Lawrence Livermore National Laboratory), *Torsten Hoefler* (ETH Zürich), *Onur Mutlu* (Carnegie Mellon University), *Keshav Pingali* (the University of Texas at Austin), *John Shalf* (Lawrence Berkeley National Laboratory)

3:00 – 3:30 Break

3:30 – 4:45 Paper Session 2: Load balancing – Chair: Alessandro Morari (PNNL)

Data-Driven Techniques to Overcome Workload Disparity (25 min)
Dan Connors (University of Colorado), Skyler Saleh (University of Colorado), Tejas Joshi (University of Colorado), and Ryan Bueter (University of Colorado)

Distributed Control: Priority Scheduling for Single Source Shortest Paths without Synchronization (25 min)
Marcin Zalewski (Indiana University), Thejaka Amila Kanewala (Indiana University), Jesun Sahariar Firoz (Indiana University), Andrew Lumsdaine (Indiana University)

When Merging and Branch Predictors Collide (25 min)
Oded Green (Georgia Institute of Technology)

4:45 – 5:35 Paper Session 3: GPU acceleration – Chair: Oreste Villa (NVIDIA)

Fast Triangle Counting on GPU (25 min)
Oded Green (ArrayFire), Pavan Yalamanchili (ArrayFire), Lluís-Miquel Munguía (Georgia Institute of Technology)

Accelerating Sparse Cholesky Factorization on GPUs (25 min)
Steven C. Rennich (NVIDIA), Darko Stosic (NVIDIA), Timothy A. Davis (Texas A&M University)