PMAM 2014 Workshop Program

08:45 – 09:00: Opening Remarks

Pavan Balaji, Minyi Guo, Zhiyi Huang

09:00 – 10:00: Keynote

Chen Ding, Professor, University of Rochester

"Program Interaction in Shared Cache: A Theory and Applications"

Abstract

On modern multicore systems, the interaction between co-run programs largely depends on cache sharing, and cache sharing depends on the locality, i.e. the active data usage, in co-run programs. This talk introduces a newly developed locality theory based on a concept called the footprint. The theory shows the composability of footprint and the conversion between footprint and other locality metrics. It enables accurate characterization and analysis of the dynamics of cache sharing.

Biography of the speaker

Chen Ding received Ph.D. from Rice University, M.S. from Michigan Tech, and B.S. from Beijing University before joining Rochester in 2000. His research received the young investigator awards from NSF and DOE. He co-founded the ACM SIGPLAN Workshop on Memory System Performance and Correctness (MSPC) and was a visiting researcher at Microsoft Research and a visiting associate professor at MIT. He is an external faculty fellow at IBM Center for Advanced Studies.

10:00- 10:30: Morning Break

10:30 – 12:00: Session 1: Runtime support

Session Chair: Bernd Burgstaller

"Work Stealing Strategies For Multi-Core Parallel Branch-and-Bound Algorithm Using Factorial Number System", Rudi Leroy, Mohand Mezmaz, Nouredine Melab and Daniel Tuyttens.

"Palirria: Accurate On-line Parallelism Estimation for Adaptive Work-Stealing", Georgios Varisteas and Mats Brorsson

“DWS: Demand-aware Work-Stealing in Multi-programmed Multi-core Architectures”, Quan Chen, Long Zheng and Minyi Guo

“Reachability Analysis of Cost-Reward Timed Automata for Energy Efficiency Scheduling”, Wei Wang, Guo Dong, Zhigang Deng, Guosun Zeng, Wei Liu and Huanliang Xiong

12:00 – 13:30: Lunch break

13:30 – 15:00: Session 2: Programming languages and models

Session Chair: Chen Ding

“Programming a Multicore Architecture without Coherency and Atomic Operations”, Jochem H. Rutgers, Marco J.G. Bekooij and Gerard J.M. Smit.

“Vectorizing Unstructured Mesh Computations for Many-core Architectures”, I. Z. Reguly, E. László, G. R. Mudalige and M. B. Giles.

“Compiling Fresh Breeze Codelets”, Jack B. Dennis.

“A Framework for Multiplatform HPC Applications”, Masayuki Ioki and Shigeru Chiba.

15:00 – 15:30: Afternoon Break

15:30 - 17:00 Session 3: Algorithms and Applications

Session Chair: TBD

“A Novel CPU-GPU Cooperative Implementation of A Parallel Two-List Algorithm for the Subset-Sum Problem”, Lanjun Wan, Kenli Li, Jing Liu and Keqin Li.

“Dynamic Partitioning-based JPEG Decompression on Heterogeneous Multicore Architectures”, Wasuwee Sodsong, Jingun Hong, Seongwook Chung, Yeongkyu Lim, Shin-Dug Kim and Bernd Burgstaller.

“Fast Longest Common Subsequence with General Integer Scoring Support on GPUs”, Adnan Ozsoy, Arun Chauhan and Martin Swany.

“Efficient Parallel Implementations of Multiple Sequence Alignment using BSP/CGM Model”, Jucele F. A. Vasconcellos, Christiane Nishibe, Nalvo F. Almeida and Edson N. Cáceres.

17:00 – 18:00: Session 4: Programming Tools

Session Chair: Xiaoyi Lu

“Autotuning Wavefront Applications for Multicore Multi-GPU Hybrid Architectures”, Siddharth Mohanty and Murray Cole.

“Reduction Operations in Parallel Loops for GPGPUs”, Rengan Xu, Xiaonan Tian, Yonghong Yan, Sunita Chandrasekaran, and Barbara Chapman.

“Self-Configuration and Self-Optimization Autonomic Skeletons using Events”, Gustavo Pabón and Ludovic Henrio

18:00 – 18:15: Closing Remarks

Pavan Balaji, Minyi Guo, Zhiyi Huang