

Table of Contents

Keynote: The Multi-core Problem as an Algorithmic Problem by Leslie Valiant of Harvard (Wed 9/28 1 – 2pm)

Large scale parallelism (Wed 9/28 2:30 – 4:05pm, 95 min)

QUARC: An Array Programming Approach to High Performance Computing 1
Diptorup Deb, Robert J. Fowler and Allan Porterfield

Utilizing Concurrency Data Access: A New Theory..... 16
Xian-He Sun

ParFuse: Parallel and Compositional Analysis of Message Passing Programs 17
Sriram Ananthakrishnan, Greg Bronevetsky, Mark Baranowski and Ganesh Gopalakrishnan

Fast Approximate Distance Queries in Unweighted Graphs using Bounded Asynchrony ... 32
Adam Fidel, Francisco Coral Sabido, Colton Riedel, Nancy Amato and Laurence Rauchwerger

Energy Avoiding Matrix Multiply 47
Kelly Livingston, Aaron Landwehr, Jose Monsalve Diaz, Stephane Zuckerman, Benoit Meister and Guang Gao

Resilience and persistence (Wed 9/28 4:15 – 5:10, 55 min)

Language Support for Reliable Memory Regions..... 62
Saurabh Hukerikar and Christian Engelmann

Harnessing Parallelism in Multicore Systems to Expedite and Improve Function Approximation 74
Aurangzeb and Rudolf Eigenmann

Adaptive Software Caching for Efficient NVRAM Data Persistence 78
Pengcheng Li and Dhruva Chakrabarti

Keynote: Towards High-Level High-Performance Software Development by P. (Saday) Sadayappan of OSU (Thu 9/29 8:30 – 9:30am)

Compiler analysis and optimization (Thu 9/29 10am – 12pm, 110 min + 10 min break)

Polyhedral Compiler Technology in Collaboration with Autotuning Important to Domain-Specific Frameworks for HPC..... 82
Mary Hall and Protonu Basu

An Extended Polyhedral Model for SPMD Programs and its use in Static Data Race Detection 86
Prasanth Chatarasi, Jun Shirako, Martin Kong and Vivek Sarkar

Polygonal Iteration Space Partitioning	101
<i>Aniket Shivam, Alexandru Nicolau, Alex V. Veidenbaum, Mario Mango Furnari and Ro Cammarota</i>	
Automatically Optimizing Stencil Computations on Many-core NUMA Architectures	116
<i>Pei-Hung Lin, Qing Yi, Daniel Quinlan, Chunhua Liao and Yongqing Yan</i>	
Tapir: Embedding Fork-Join Parallelism into LLVM's Intermediate Representation	131
<i>Tao Schardl</i>	
Formalizing Structured Control Flow Graphs	132
<i>Amit Sabne, Putt Sakdhnagool and Rudolf Eigenmann</i>	
<hr/> Lunch (provided, 12 – 1pm) <hr/>	
<hr/> Dynamic computation and languages (Thu 9/29 1 – 2:20pm, 80 min) <hr/>	
Automatic Vectorization for MATLAB	136
<i>Hanfeng Chen, Alexander Krolik, Erick Lavoie and Laurie Hendren</i>	
Analyzing Parallel Programming Models for Magnetic Resonance Imaging	151
<i>Forest Danford, Eric Welch, Julio Cárdenas-Rodríguez and Michelle Strout</i>	
The Importance of Fine-Grain Synchronization for Many-Core Systems	166
<i>Tongsheng Geng, Stéphane Zuckerman, Jose Monsalve, Alfredo Goldman, Sami Habib, Jean-Luc Gaudiot and Guang R. Gao</i>	
Optimizing LOBPCG: Sparse Matrix Loop and Data Transformations in Action	181
<i>Khalid Ahmad, Anand Venkat and Mary Hall</i>	
<hr/> Julia compiler tutorial and QA by Keno Fischer (Thu 9/29 2:30 – 3:30pm) <hr/>	
<hr/> Panel: Compilation for dynamic parallel languages, by Ayon Basumallik of MathWorks, Keno Fischer of Julia, and Chu-Cheow Lim of Qualcomm (Thu 9/29 4 – 5pm) <hr/>	
<hr/> Excursion: Kodak theatre Break of Reality concert (8pm to 10pm) or Erie Canal cruise/dinner (6:30pm to 9:30pm), leaving from the hotel half hour before start time (transportation will be provided from and back to the hotel) <hr/>	
<hr/> Keynote: Parallel Computation Models and Systems, Dataflow, Coelets, and Beyond by Guang R. Gao of U. Delaware (Fri 9/30 8:30 – 9:30am) <hr/>	
<hr/> GPUs and private memory (Fri 9/30 10 – 11:20am, 80 min) <hr/>	
LightHouse: An Automatic Code Generator for Graph Algorithms on GPUs	196
<i>Shashidhar G and Rupesh Nasre.</i>	
Locality-aware Task-parallel Execution on GPUs	211
<i>Jad Hbeika and Milind Kulkarni</i>	
Automatic Copying of Pointer-Based Data Structures	226
<i>Tong Chen, Zehra Sura and Hyojin Sung</i>	

Automatic Local Memory Management for Multicores Having Global Address Space	241
<i>Kouhei Yamamoto, Tomoya Shirakawa, Yoshitake Oki, Akimasa Yoshida, Keiji Kimura and Hironori Kasahara</i>	
<hr/> Run-time and performance analysis (Fri 9/30 11:30 – 12:30, 60 min) <hr/>	
Mapping Medley: Adaptive Parallelism Mapping with Varying Optimization Goals	256
<i>Murali Emani</i>	
The Contention Avoiding Concurrent Priority Queue	271
<i>Kjell Winblad and Konstantinos Sagonas</i>	
Evaluating Performance of Task and Data Coarsening in Concurrent Collections	286
<i>Chenyang Liu and Milind Kulkarni</i>	