

Contents

Large Scale Parallelism

QUARC: An Array Programming Approach to High Performance Computing	3
<i>Diptorup Deb, Robert J. Fowler, and Allan Porterfield</i>	
Utilizing Concurrency: A New Theory for Memory Wall	18
<i>Xian-He Sun and Yu-Hang Liu</i>	
ParFuse: Parallel and Compositional Analysis of Message Passing Programs	24
<i>Sriram Aananthakrishnan, Greg Bronevetsky, Mark Baranowski, and Ganesh Gopalakrishnan</i>	
Fast Approximate Distance Queries in Unweighted Graphs Using Bounded Asynchrony	40
<i>Adam Fidel, Francisco Coral Sabido, Colton Riedel, Nancy M. Amato, and Lawrence Rauchwerger</i>	
Energy Avoiding Matrix Multiply	55
<i>Kelly Livingston, Aaron Landwehr, José Monsalve, Stéphane Zuckerman, Benoît Meister, and Guang R. Gao</i>	

Resilience and Persistence

Language Support for Reliable Memory Regions	73
<i>Saurabh Hukerikar and Christian Engelmann</i>	
Harnessing Parallelism in Multicore Systems to Expedite and Improve Function Approximation	88
<i>Aurangzeb and Rudolf Eigenmann</i>	
Adaptive Software Caching for Efficient NVRAM Data Persistence	93
<i>Pengcheng Li and Dhruva R. Chakrabarti</i>	

Compiler Analysis and Optimization

Polyhedral Compiler Technology in Collaboration with Autotuning Important to Domain-Specific Frameworks for HPC	101
<i>Mary Hall and Protonu Basu</i>	

An Extended Polyhedral Model for SPMD Programs and Its Use in Static Data Race Detection	106
<i>Prasanth Chatarasi, Jun Shirako, Martin Kong, and Vivek Sarkar</i>	
Polygonal Iteration Space Partitioning	121
<i>Aniket Shivam, Alexandru Nicolau, Alexander V. Veidenbaum, Mario Mango Furnari, and Rosario Cammarota</i>	
Automatically Optimizing Stencil Computations on Many-Core NUMA Architectures	137
<i>Pei-Hung Lin, Qing Yi, Daniel Quinlan, Chunhua Liao, and Yongqing Yan</i>	
Formalizing Structured Control Flow Graphs	153
<i>Amit Sabne, Putt Sakdhnagool, and Rudolf Eigenmann</i>	
Dynamic Computation and Languages	
Automatic Vectorization for MATLAB	171
<i>Hanfeng Chen, Alexander Krolik, Erick Lavoie, and Laurie Hendren</i>	
Analyzing Parallel Programming Models for Magnetic Resonance Imaging.	188
<i>Forest Danford, Eric Welch, Julio Cárdenas-Rodríguez, and Michelle Mills Strout</i>	
The Importance of Efficient Fine-Grain Synchronization for Many-Core Systems	203
<i>Tongsheng Geng, Stéphane Zuckerman, José Monsalve, Alfredo Goldman, Sami Habib, Jean-Luc Gaudiot, and Guang R. Gao</i>	
Optimizing LOBPCG: Sparse Matrix Loop and Data Transformations in Action	218
<i>Khalid Ahmad, Anand Venkat, and Mary Hall</i>	
GPUs and Private Memory	
LightHouse: An Automatic Code Generator for Graph Algorithms on GPUs	235
<i>G. Shashidhar and Rupesh Nasre</i>	
Locality-Aware Task-Parallel Execution on GPUs	250
<i>Jad Hbeika and Milind Kulkarni</i>	
Automatic Copying of Pointer-Based Data Structures	265
<i>Tong Chen, Zehra Sura, and Hyojin Sung</i>	

Automatic Local Memory Management for Multicores Having Global Address Space	282
<i>Kouhei Yamamoto, Tomoya Shirakawa, Yoshitake Oki, Akimasa Yoshida, Keiji Kimura, and Hironori Kasahara</i>	
Run-time and Performance Analysis	
Mapping Medley: Adaptive Parallelism Mapping with Varying Optimization Goals	299
<i>Murali Krishna Emami</i>	
The Contention Avoiding Concurrent Priority Queue.	314
<i>Konstantinos Sagonas and Kjell Winblad</i>	
Evaluating Performance of Task and Data Coarsening in Concurrent Collections.	331
<i>Chenyang Liu and Milind Kulkarni</i>	
Author Index	347

Languages and Compilers for Parallel Computing
29th International Workshop, LCPC 2016, Rochester,
NY, USA, September 28-30, 2016, Revised Papers
Ding, C.; Criswell, J.; Wu, P. (Eds.)
2017, XI, 348 p. 137 illus., Softcover
ISBN: 978-3-319-52708-6