

Contents

Programming Models

Hierarchical Computation in the SPMD Programming Model	3
<i>Amir Kamil and Katherine Yelick</i>	
Porting Applications with OpenMP Using Similarity Analysis	20
<i>Wei Ding, Oscar Hernandez, Tony Curtis, and Barbara Chapman</i>	

Tasks

Task-Aware Optimization of Dynamic Fractional Permissions.	39
<i>Christoph M. Angerer</i>	
Near Optimal Work-Stealing Tree Scheduler for Highly Irregular Data-Parallel Workloads.	55
<i>Aleksandar Prokopec and Martin Odersky</i>	
OpenCL Task Partitioning in the Presence of GPU Contention	87
<i>Dominik Grewe, Zheng Wang, and Michael F.P. O’Boyle</i>	

Heterogeneous Computing

Compiling a High-Level Directive-Based Programming Model for GPGPUs . . .	105
<i>Xiaonan Tian, Rengan Xu, Yonghong Yan, Zhifeng Yun, Sunita Chandrasekaran, and Barbara Chapman</i>	
Separate Compilation in a Language-Integrated Heterogeneous Environment. . .	121
<i>Mike Murphy, Jaydeep Marathe, Girish Bharambe, Sean Lee, and Vinod Grover</i>	
Parametric GPU Code Generation for Affine Loop Programs	136
<i>Athanasios Konstantinidis, Paul H.J. Kelly, J. Ramanujam, and P. Sadayappan</i>	

Power

OSCAR Compiler Controlled Multicore Power Reduction on Android Platform	155
<i>Hideo Yamamoto, Tomohiro Hirano, Kohei Muto, Hiroki Mikami, Takashi Goto, Dominic Hillenbrand, Moriyuki Takamura, Keiji Kimura, and Hironori Kasahara</i>	

Folklore Confirmed: Compiling for Speed = Compiling for Energy	169
<i>Tomofumi Yuki and Sanjay Rajopadhye</i>	

Debugging

Effectively Recognize Ad hoc Synchronizations with Static Analysis.	187
<i>Le Yin</i>	
AntSM: Efficient Debugging for Shared Memory Parallel Programs	202
<i>Jae-Woo Lee and Samuel P. Midkiff</i>	
DRIFT: Decoupled CompileR-Based Instruction-Level Fault-Tolerance	217
<i>Konstantina Mitropoulou, Vasileios Porpodas, and Marcelo Cintra</i>	

Algorithms

Optimizing the LU Factorization for Energy Efficiency on a Many-Core Architecture	237
<i>Elkin Garcia, Jaime Arteaga, Robert Pavel, and Guang R. Gao</i>	
An Input-Adaptive Algorithm for High Performance Sparse Fast Fourier Transform.	252
<i>Shuo Chen and Xiaoming Li</i>	

Caches

Aligned Scheduling: Cache-Efficient Instruction Scheduling for VLIW Processors.	275
<i>Vasileios Porpodas and Marcelo Cintra</i>	
Compile Time Modeling of Off-Chip Memory Bandwidth for Parallel Loops . . .	292
<i>Munara Tolubaeva, Yonghong Yan, and Barbara Chapman</i>	
Compiler Optimizations for Non-contiguous Remote Data Movement	307
<i>Timo Schneider, Robert Gerstenberger, and Torsten Hoefler</i>	

Transactional Memory

Combining Lock Inference with Lock-Based Software Transactional Memory . . .	325
<i>Stefan Kempf, Ronald Veldema, and Michael Philippsen</i>	
Speculative Execution of Parallel Programs with Precise Exception Semantics on GPUs	342
<i>Akihiro Hayashi, Max Grossman, Jisheng Zhao, Jun Shirako, and Vivek Sarkar</i>	

Author Index	357
------------------------	-----

Languages and Compilers for Parallel Computing

26th International Workshop, LCPC 2013, San Jose, CA,
USA, September 25--27, 2013. Revised Selected Papers

Caşcaval, C.; Montesinos, P. (Eds.)

2014, XXIV, 357 p. 160 illus., Softcover

ISBN: 978-3-319-09966-8