

Table of Contents

Automatic Scaling of OpenMP Beyond Shared Memory	1
<i>Okwan Kwon, Fahed Jubair, Seung-Jai Min, Hansang Bae, Rudolf Eigenmann, and Samuel P. Midkiff</i>	
A Methodology for Fine-Grained Parallelism in JavaScript Applications	16
<i>Jeffrey Fifield and Dirk Grunwald</i>	
Evaluation of Power Consumption at Execution of Multiple Automatically Parallelized and Power Controlled Media Applications on the RP2 Low-Power Multicore	31
<i>Hiroki Mikami, Shumpei Kitaki, Masayoshi Mase, Akihiro Hayashi, Mamoru Shimaoka, Keiji Kimura, Masato Edahiro, and Hironori Kasahara</i>	
Double Inspection for Run-Time Loop Parallelization	46
<i>Michael Philippsen, Nikolai Tillmann, and Daniel Brinkers</i>	
A Hybrid Approach to Proving Memory Reference Monotonicity	61
<i>Cosmin E. Oancea and Lawrence Rauchwerger</i>	
OpenCL as a Programming Model for GPU Clusters	76
<i>Jungwon Kim, Sangmin Seo, Jun Lee, Jeongho Nah, Gangwon Jo, and Jaejin Lee</i>	
CellCilk: Extending Cilk for Heterogeneous Multicore Platforms	91
<i>Tobias Werth, Silvia Schreier, and Michael Philippsen</i>	
OPELL and PM: A Case Study on Porting Shared Memory Programming Models to Accelerators Architectures	106
<i>Joseph B. Manzano, Ge Gan, Juergen Ributzka, Sunil Shrestha, and Guang R. Gao</i>	
Optimizing the Concurrent Execution of Locks and Transactions	124
<i>Justin E. Gottschlich and JaeWoong Chung</i>	
A Study of the Usefulness of Producer/Consumer Synchronization	141
<i>Hao Lin, Hansang Bae, Samuel P. Midkiff, Rudolf Eigenmann, and Soohong P. Kim</i>	
Lock-Free Resizeable Concurrent Tries	156
<i>Aleksandar Prokopec, Phil Bagwell, and Martin Odersky</i>	

Fine-Grained Treatment to Synchronizations in GPU-to-CPU Translation	171
<i>Ziyu Guo and Xipeng Shen</i>	
A Mutable Hardware Abstraction to Replace Threads	185
<i>Sean Halle and Albert Cohen</i>	
Dynamic Task Parallelism with a GPU Work-Stealing Runtime System	203
<i>Sanjay Chatterjee, Max Grossman, Alina Sbîrlea, and Vivek Sarkar</i>	
A Code Merging Optimization Technique for GPU	218
<i>Ryan Taylor and Xiaoming Li</i>	
Static Compilation Analysis for Host-Accelerator Communication Optimization	237
<i>Mehdi Amini, Fabien Coelho, François Irigoín, and Ronan Keryell</i>	
Scheduling Support for Communicating Parallel Tasks	252
<i>Jörg Dümmler, Thomas Rauber, and Gudula Rünger</i>	
Polytasks: A Compressed Task Representation for HPC Runtimes	268
<i>Daniel Orozco, Elkin Garcia, Robert Pavel, Rishi Khan, and Guang R. Gao</i>	
Detecting False Sharing in OpenMP Applications Using the DARWIN Framework	283
<i>Besar Wicaksono, Munara Tolubaeva, and Barbara Chapman</i>	
Author Index	299