2014 IEEE 26th International Symposium on Computer Architecture and High Performance Computing

(SBAC-PAD 2014)

Paris, France
22 - 24 October 2014

IEEE Catalog Number: CFP14307-POD
ISBN: 978-1-4799-6906-7

2014 IEEE 26th International Symposium on Computer Architecture and High Performance Computing

SBAC-PAD 2014

Table of Contents

Message from the General Chairs x

Message from the Program and Track Chairs xi

Committee Lists xiv

Keynotes xviii

Shared-Memory and Multi-core Systems

A Case Study of Hybrid Dataflow and Shared-Memory Programming Models: Dependency-Based Parallel Game Engine 1
Vladimir Gajinov, Igor Erić, Saša Stojanović, Veljko Milutinović, Osman Unsal, Eduard Ayguadé, and Adrián Cristal

Wait-Free Global Virtual Time Computation in Shared Memory TimeWarp Systems 9
Alessandro Pellegrini and Francesco Quaglia

Compact Hash Tables for High-Performance Traffic Classification on Multi-core Processors 17
Yun R. Qu and Viktor K. Prasanna

Performance Optimization

Flying Memcache: Lessons Learned from Different Acceleration Strategies 25
Dimitris Deyannis, Lazaros Koromilas, Giorgos Vasiliadis, Elias Athanasopoulos, and Sotiris Ioannidis

Improving an MPI Application-Level Migration Approach through Checkpoint File Splitting 33
Mónica Rodríguez, Iván Cores, Patricia González, and María J. Martín

HPCG: Preliminary Evaluation and Optimization on Tianhe-2 CPU-only Nodes 41
Cheng Chen, Yunfei Du, Hao Jiang, Ke Zuo, and Canqun Yang

Leveraging Optimization Methods for Dynamically Assisted Control-Flow Integrity Mechanisms 49
João Moreira, Lucas Teixeira, Edson Borin, and Sandro Rigo

Energy Aware Computing

Energy Efficient Seismic Wave Propagation Simulation on a Low-Power Manycore Processor 57
Márcio Castro, Fabrice Dupros, Emilio Francesquini, Jean-François Méhautk, and Philippe O.A. Navaux

Performance-Aware Task Management and Frequency Scaling in Embedded Systems 65
Francisco Gaspar, Aleksandar Ilic, Pedro Tomás, and Leonel Sousa

Analyzing Performance Improvements and Energy Savings in Infiniband Architecture using Network Compression 73
Branimir Dickov, Miquel Pericàs, Paul M. Carpenter, Nacho Navarro, and Eduard Ayguadé

Hardware Accelerators I

Bit-Parallel Approximate Pattern Matching on the Xeon Phi Coprocessor 81
Tuan Tu Tran, Simon Schindel, Yongchao Liu, and Bertil Schmidt

Efficient Execution of Microscopy Image Analysis on CPU, GPU, and MIC Equipped Cluster Systems 89
G. Andrade, R. Ferreira, George Teodoro, Leonardo Rocha, Joel H. Saltz, and Tahsin Kurc

High-Performance Traffic Classification on GPU 97
Shijie Zhou, Prashant Rao Nittoor, and Viktor K. Prasanna

Accelerating Curvature Estimate in 3D Seismic Data Using GPGPU 105
Leonardo Martins, Marco Aurélio Gonçalves da Silva, Marcelo Arruda, Joner Duarte Jr, Pedro Mário Silva, Roberto Beauclair Seixas, and Marcelo Gattass

Leveraging OmpSs to Exploit Hardware Accelerators 112
Florentino Sainz, Sergi Mateo, Vicenç Beltran, Jose L. Bosque, Xavier Martorell, and Eduard Ayguadé

Transactional Memory

Improving Signature Behavior by Irrevocability in Transactional Memory Systems 120
Ricardo Quislant, Eladio Gutierrez, Emilio L. Zapata, and Oscar Plata

Scalability Analysis of Signatures in Transactional Memory Systems 128
Ricardo Quislant, Eladio Gutierrez, and Oscar Plata

Profiling Patterns of Bit Flipping for Software Transactional Memories 136
Felipe L. Teixeira, Maurício L. Pilla, André R. Du Bois, and Daniel Mossé

Multi-dimensional Evaluation of Haswell's Transactional Memory Performance 144
Marcio Machado Pereira, Matthew Gaudet, José Nelson Amaral, and Guido Araújo

DeTrans: Deterministic and Parallel execution of Transactions 152
Vesna Smiljković, Srđan Stipić, Christof Fetzer, Osman Ünsal, Adrián Cristal, and Mateo Valero

Hardware Accelerators II

Design Space Exploration of Memory Model for Heterogeneous Computing 160
Jieun Lim and Hyesoon Kim

Runtime Support for Adaptive Spatial Partitioning and Inter-Kernel Communication on GPUs 168
Yash Ukidave, Charu Kalra, David Kaeli, Perhaad Mistry, and Dana Schaa

Automatic Generation of Custom Parallel Processors for Morphological Image Processing 176
Emerson Carlos Pedrino and Marcio Merino Fernandes

ASIC Design of Shared Vector Accelerators for Multicore Processors 182
Spiridon F. Beldianu and Sotirios G. Ziavras

Efficient Scan Operator Methods on a GPU 190
Adrián P. Diéguez, Margarita Amor, and Ramón Doallo

Memory Hierarchy

Optimizing Memory Locality Using a Locality-Aware Page Table 198
Eduardo H.M. Cruz, Matthias Diener, Marco A.Z. Alves, Laércio L. Pilla, and Philippe O.A. Navaux

HiTS: A High Throughput Memory Scheduling Scheme to Mitigate Denial-of-Service Attacks in Multi-core Systems 206
Mansour Shafaei and Yunsi Fei

Hardware/Software Helper Thread Prefetching on Heterogeneous Many Cores 214
Bharath Narasimha Swamy, Alain Ketterlin, and André Seznec

Chip-Level Optimizations

Profiling and Reducing Micro-Architecture Bottlenecks at the Hardware Level 222
Francis B. Moreira, Marco A.Z. Alves, and Israel Koren

Design and Evaluation of Hierarchical Rings with Deflection Routing 230
Rachata Ausavarungnirun, Chris Fallin, Xiangyao Yu, Kevin Kai-Wei Chang, Greg Nazario, Reetuparna Das, Gabriel H. Loh, and Onur Mutlu

Block Disabling Characterization and Improvements in CMPs Operating at Ultra-low Voltages 238
Alexandra Ferrerón, Darío Suárez-Gracia, Jesús Alastruey-Benedé, Teresa Monreal, and Víctor Viñals

Fine-Grained Parallelism

Evaluation of a Feature Tracking Vision Application on a Heterogeneous Chip 246
Rubén Gran, August Shi, Ehsan Totoni, and María J. Garzarán

Predicting GPU Performance from CPU Runs Using Machine Learning 254
Ioana Baldini, Stephen J. Fink, and Erik Altman

Leveraging Task-Parallelism with OmpSs in ILUPACK's Preconditioned CG Method 262
J.I. Aliaga, R.M. Badía, M. Barreda, M. Bollhöfer, and E.S. Quintana-Ortí

Reducing Compiler-Inserted Instrumentation in Unified-Parallel-C Code Generation 270
Michail Alvanosl, José Nelson Amaral, Ettore Tiotto, Montse Farreras, and Xavier Martorell

Lock-Free GaussSieve for Linear Speedups in Parallel High Performance SVP Calculation 278
Artur Mariano, Shahar Timnat, and Christian Bischof

Grid and Cloud Computing

Wide Area BonjourGrid as a Data Desktop Grid: Modeling and Implementation on Top of Redis 286
Walid Saad, Leila Abidi, Heithem Abbes, Christophe Cérin, and Mohamed Jemni

An Anomaly Detection Approach for Scale-Out Storage Systems 294
Guthemberg Silvestre, Carla Sauvanaud, Mohamed Kaâniche, and Karama Kanoun

Analyzing Real Cluster Data for Formulating Allocation Algorithms in Cloud Platforms 302
Olivier Beaumont, Lionel Eyraud-Dubois, and Juan-Angel Lorenzo-del-Castillo

Modeling the Impact of Workload on Cloud Resource Scaling 310
Anshul Gandhi, Parijat Dube, Alexei Karve, Andrzej Kochut, and Li Zhang

Parallel and Distributed Systems

SLURM Support for Remote GPU Virtualization: Implementation and Performance Study 318
Sergio Iserte, Adrián Castelló, Rafael Mayo, Enrique S. Quintana-Ortí, Federico Silla, Jose Duato, Carlos Reaño, and Javier Prades

Cross-Layer Self-Adaptive/Self-Aware System Software for Exascale Systems 326
Roberto Gioiosa, Gokcen Kestor, Darren J. Kerbyson, and Adolfy Hoisie

Cloud-Based OpenMP Parallelization Using a MapReduce Runtime 334
Rodolfo Wottrich, Rodolfo Azevedo, and Guido Araújo

Author Index 342