

# Contents

## Invited Papers

Resampling with Feedback — A New Paradigm of Using Workload Data for Performance Evaluation . . . . .	3
<i>Dror G. Feitelson</i>	
Scheduling DAGs Opportunistically: The Dream and the Reality Circa 2016 . . . . .	22
<i>Arnold L. Rosenberg</i>	

## Support Tools and Environments

Synchronization Debugging of Hybrid Parallel Programs . . . . .	37
<i>Olaf Krzikalla, Ralph Müller-Pfefferkorn, and Wolfgang E. Nagel</i>	
Nasty-MPI: Debugging Synchronization Errors in MPI-3 One-Sided Applications . . . . .	51
<i>Roger Kowalewski and Karl Fürlinger</i>	
Automatic Benchmark Profiling Through Advanced Trace Analysis . . . . .	63
<i>Alexis Martin and Vania Marangozova-Martin</i>	

## Performance and Power Modeling, Prediction and Evaluation

Addressing Materials Science Challenges Using GPU-accelerated POWER8 Nodes . . . . .	77
<i>Paul F. Baumeister, Marcel Bornemann, Markus Bühler, Thorsten Hater, Benjamin Krill, Dirk Pleiter, and Rudolf Zeller</i>	
Performance Prediction and Ranking of SpMV Kernels on GPU Architectures . . . . .	90
<i>Christoph Lehnert, Rudolf Berrendorf, Jan P. Ecker, and Florian Mannuss</i>	
The Impact of Voltage-Frequency Scaling for the Matrix-Vector Product on the IBM POWER8 . . . . .	103
<i>Sandra Catalán, A. Cristiano I. Malossi, Costas Bekas, and Enrique S. Quintana-Ortí</i>	
Power Consumption Modeling and Prediction in a Hybrid CPU-GPU-MIC Supercomputer . . . . .	117
<i>Alina Sîrbu and Ozalp Babaoglu</i>	

## Scheduling and Load Balancing

Controlling and Assessing Correlations of Cost Matrices in Heterogeneous Scheduling . . . . .	133
<i>Louis-Claude Canon, Pierre-Cyrille Héam, and Laurent Philippe</i>	
Penalized Graph Partitioning for Static and Dynamic Load Balancing . . . . .	146
<i>Tim Kiefer, Dirk Habich, and Wolfgang Lehner</i>	
Non-preemptive Scheduling with Setup Times: A PTAS . . . . .	159
<i>Klaus Jansen and Felix Land</i>	
Cuboid Partitioning for Parallel Matrix Multiplication on Heterogeneous Platforms . . . . .	171
<i>Olivier Beaumont, Lionel Eyraud-Dubois, and Thomas Lambert</i>	
HeSP: A Simulation Framework for Solving the Task Scheduling-Partitioning Problem on Heterogeneous Architectures . . . . .	183
<i>Antón Rey, Francisco D. Igual, and Manuel Prieto-Matías</i>	
FPT Approximation Algorithm for Scheduling with Memory Constraints . . . .	196
<i>Eric Angel, Cédric Chevalier, Franck Ledoux, Sébastien Morais, and Damien Regnault</i>	
Scheduling MapReduce Jobs Under Multi-round Precedences . . . . .	209
<i>D. Fotakis, I. Milis, O. Papadigenopoulos, V. Vassalos, and G. Zois</i>	

## High Performance Architectures and Compilers

Code Bones: Fast and Flexible Code Generation for Dynamic and Speculative Polyhedral Optimization . . . . .	225
<i>Juan Manuel Martinez Caamaño, Willy Wolff, and Philippe Clauss</i>	
Piecewise Holistic Autotuning of Compiler and Runtime Parameters . . . . .	238
<i>Mihail Popov, Chadi Akel, William Jalby, and Pablo de Oliveira Castro</i>	
Insights into the Fallback Path of Best-Effort Hardware Transactional Memory Systems . . . . .	251
<i>Ricardo Quislan, Eladio Gutierrez, Emilio L. Zapata, and Oscar Plata</i>	
Portable SIMD Performance with OpenMP* 4.x Compiler Directives . . . . .	264
<i>Florian Wende, Matthias Noack, Thomas Steinke, Michael Klemm, Chris J. Newburn, and Georg Zitzlsberger</i>	

## Parallel and Distributed Data Management and Analytics

Lightweight Multi-language Bindings for Apache Spark. . . . .	281
<i>Luca Salucci, Daniele Bonetta, and Walter Binder</i>	

Toward a General I/O Arbitration Framework for netCDF Based Big Data Processing . . . . .	293
<i>Jianwei Liao, Balazs Gerofi, Guo-Yuan Lien, Seiya Nishizawa, Takemasa Miyoshi, Hirofumi Tomita, and Yutaka Ishikawa</i>	
High Performance Parallel Summed-Area Table Kernels for Multi-core and Many-core Systems . . . . .	306
<i>Angelos Papatriantafyllou and Dimitris Sacharidis</i>	
GraphIn: An Online High Performance Incremental Graph Processing Framework . . . . .	319
<i>Dipanjana Sengupta, Narayanan Sundaram, Xia Zhu, Theodore L. Willke, Jeffrey Young, Matthew Wolf, and Karsten Schwan</i>	
Efficient Large Outer Joins over MapReduce . . . . .	334
<i>Long Cheng and Spyros Kotoulas</i>	

## Cluster and Cloud Computing

Slurm-V: Extending Slurm for Building Efficient HPC Cloud with SR-IOV and IVShmem . . . . .	349
<i>Jie Zhang, Xiaoyi Lu, Sourav Chakraborty, and Dhabaleswar K. (DK) Panda</i>	
An Autonomic Parallel Strategy for the Projection of Ecological Niche Models in Heterogeneous Computational Environments . . . . .	363
<i>Fernanda G.O. Passos and Vinod E.F. Rebello</i>	
Towards Network-Aware Service Placement in Community Network Micro-Clouds . . . . .	376
<i>Mennan Selimi, Davide Vega, Felix Freitag, and Luís Veiga</i>	
Heating as a Cloud-Service, A Position Paper (Industrial Presentation) . . . . .	389
<i>Yanik Ngoko</i>	

## Distributed Systems and Algorithms

Design and Verification of Distributed Phasers . . . . .	405
<i>Karthik Murthy, Sri Raj Paul, Kuldeep S. Meel, Tiago Cogumbreiro, and John Mellor-Crummey</i>	
Exploring Partial Replication to Improve Lightweight Silent Data Corruption Detection for HPC Applications . . . . .	419
<i>Eduardo Berrocal, Leonardo Bautista-Gomez, Sheng Di, Zhiling Lan, and Franck Cappello</i>	

## Parallel and Distributed Programming, Interfaces, Language

Automatic Verification of Self-consistent MPI Performance Guidelines . . . . .	433
<i>Sascha Hunold, Alexandra Carpen-Amarie, Felix Donatus Lübke, and Jesper Larsson Träff</i>	
ParallelME: A Parallel Mobile Engine to Explore Heterogeneity in Mobile Computing Architectures . . . . .	447
<i>Guilherme Andrade, Wilson de Carvalho, Renato Utsch, Pedro Caldeira, Alberto Alburquerque, Fabricio Ferracioli, Leonardo Rocha, Michael Frank, Dorgival Guedes, and Renato Ferreira</i>	
CBPQ: High Performance Lock-Free Priority Queue . . . . .	460
<i>Anastasia Braginsky, Nachshon Cohen, and Erez Petrank</i>	

## Multicore and Manycore Parallelism

Redesigning Triangular Dense Matrix Computations on GPUs . . . . .	477
<i>Ali Charara, Hatem Ltaief, and David Keyes</i>	
A Sharing-Aware Memory Management Unit for Online Mapping in Multi-core Architectures . . . . .	490
<i>Eduardo H.M. Cruz, Matthias Diener, Laércio L. Pilla, and Philippe O.A. Navaux</i>	
GreenBST: Energy-Efficient Concurrent Search Tree . . . . .	502
<i>Ibrahim Umar, Otto Anshus, and Phuong Ha</i>	
HAP: A Heterogeneity-Conscious Runtime System for Adaptive Pipeline Parallelism . . . . .	518
<i>Jinsu Park and Woongki Baek</i>	
Using Data Dependencies to Improve Task-Based Scheduling Strategies on NUMA Architectures . . . . .	531
<i>Philippe Virouleau, François Broquedis, Thierry Gautier, and Fabrice Rastello</i>	
Multicore vs Manycore: The Energy Cost of Concurrency . . . . .	545
<i>Martin Groen and Vincent Gramoli</i>	

## Theory and Algorithms for Parallel Computation and Networking

Work-Efficient Parallel Union-Find with Applications to Incremental Graph Connectivity . . . . .	561
<i>Natcha Simsiri, Kanat Tangwongsan, Srikanta Tirthapura, and Kun-Lung Wu</i>	

An Efficient Cache-oblivious Parallel Viterbi Algorithm . . . . .	574
<i>Rezaul Chowdhury, Pramod Ganapathi, Vivek Pradhan, Jesmin Jahan Tithi, and Yunpeng Xiao</i>	
Gradual Stabilization Under $\tau$ -Dynamics . . . . .	588
<i>Karine Altisen, Stéphane Devismes, Anaïs Durand, and Franck Petit</i>	
<b>Parallel Numerical Methods and Applications</b>	
High Performance Polar Decomposition on Distributed Memory Systems. . . .	605
<i>Dalal Sukkari, Hatem Ltaief, and David Keyes</i>	
A Synchronization-Free Algorithm for Parallel Sparse Triangular Solves . . . .	617
<i>Weifeng Liu, Ang Li, Jonathan Hogg, Iain S. Duff, and Brian Vinter</i>	
Exploiting Task-Parallelism in Message-Passing Sparse Linear System Solvers Using OmpSs . . . . .	631
<i>José I. Aliaga, María Barreda, Matthias Bollhöfer, and Enrique S. Quintana-Ortí</i>	
Lightweight and Accurate Silent Data Corruption Detection in Ordinary Differential Equation Solvers . . . . .	644
<i>Pierre-Louis Guhur, Hong Zhang, Tom Peterka, Emil Constantinescu, and Franck Cappello</i>	
<b>Accelerator Computing</b>	
High-Performance Matrix-Matrix Multiplications of Very Small Matrices. . . .	659
<i>Ian Masliah, Ahmad Abdelfattah, A. Haidar, S. Tomov, Marc Baboulin, J. Falcou, and J. Dongarra</i>	
Effective Minimally-Invasive GPU Acceleration of Distributed Sparse Matrix Factorization . . . . .	672
<i>Anshul Gupta, Natalia Gimelshein, Seid Koric, and Steven Rennich</i>	
Automatic OpenCL Task Adaptation for Heterogeneous Architectures . . . . .	684
<i>Pierre Huchant, Marie-Christine Counilh, and Denis Barthou</i>	
<b>Author Index</b> . . . . .	697

Euro-Par 2016: Parallel Processing  
22nd International Conference on Parallel and  
Distributed Computing, Grenoble, France, August  
24-26, 2016, Proceedings  
Dutot, P.-F.; Trystram, D. (Eds.)  
2016, XXIX, 699 p. 236 illus., Softcover  
ISBN: 978-3-319-43658-6