

Contents

Invited Paper

Computing Just What You Need: Online Data Analysis and Reduction at Extreme Scales	3
<i>Ian Foster, Mark Ainsworth, Bryce Allen, Julie Bessac, Franck Cappello, Jong Youl Choi, Emil Constantinescu, Philip E. Davis, Sheng Di, Wendy Di, Hanqi Guo, Scott Klasky, Kerstin Kleese Van Dam, Tahsin Kurc, Qing Liu, Abid Malik, Kshitij Mehta, Klaus Mueller, Todd Munson, George Ostouchov, Manish Parashar, Tom Peterka, Line Pouchard, Dingwen Tao, Ozan Tugluk, Stefan Wild, Matthew Wolf, Justin M. Wozniak, Wei Xu, and Shinjae Yoo</i>	

Support Tools and Environments

Scaling Energy Adaptive Applications for Sustainable Profitability	23
<i>Fabien Hermenier, Giuliani Giovanni, Andre Milani, and Sophie Demassey</i>	
Off-Road Performance Modeling – How to Deal with Segmented Data	36
<i>M. Kashif Ilyas, Alexandru Calotoiu, and Felix Wolf</i>	
Online Dynamic Monitoring of MPI Communications	49
<i>George Bosilca, Clément Foyer, Emmanuel Jeannot, Guillaume Mercier, and Guillaume Papauré</i>	

Performance and Power Modeling, Prediction and Evaluation

Micro-benchmarking MPI Neighborhood Collective Operations	65
<i>Felix Donatus Lübke</i>	
Performance Characterization of De Novo Genome Assembly on Leading Parallel Systems	79
<i>Marquita Ellis, Evangelos Georganas, Rob Egan, Steven Hofmeyr, Aydin Buluç, Brandon Cook, Leonid Oliker, and Katherine Yelick</i>	
NVIDIA Jetson Platform Characterization	92
<i>Hassan Halawa, Hazem A. Abdelhafez, Andrew Boktor, and Matei Ripeanu</i>	

Following the Blind Seer – Creating Better Performance Models Using Less Information	106
<i>Patrick Reisert, Alexandru Calotoiu, Sergei Shudler, and Felix Wolf</i>	
An Accurate Simulator of Cache-Line Conflicts to Exploit the Underlying Cache Performance	119
<i>Yukinori Sato and Toshio Endo</i>	
Shutdown Policies with Power Capping for Large Scale Computing Systems.	134
<i>Anne Benoit, Laurent Lefèvre, Anne-Cécile Orgerie, and Issam Raïs</i>	
Scheduling and Load Balancing	
Partitioning Strategy Selection for In-Memory Graph Pattern Matching on Multiprocessor Systems.	149
<i>Alexander Krause, Thomas Kissinger, Dirk Habich, Hannes Voigt, and Wolfgang Lehner</i>	
Efficient Dynamic Pinning of Parallelized Applications by Reinforcement Learning with Applications.	164
<i>Georgios C. Chasparis, Michael Rossbory, and Vladimir Janjic</i>	
Accelerating by Idling: How Speculative Delays Improve Performance of Message-Oriented Systems	177
<i>Aleksandar Prokopec</i>	
Using Simulation to Evaluate and Tune the Performance of Dynamic Load Balancing of an Over-Decomposed Geophysics Application	192
<i>Rafael Keller Tesser, Lucas Mello Schnorr, Arnaud Legrand, Fabrice Dupros, and Philippe Olivier Alexandre Navaux</i>	
Optimizing Egalitarian Performance in the Side-Effects Model of Colocation for Data Center Resource Management	206
<i>Fanny Pascual and Krzysztof Rzađca</i>	
Generic Algorithms for Scheduling Applications on Hybrid Multi-core Machines	220
<i>Marcos Amaris, Giorgio Lucarelli, Clément Mommessin, and Denis Trystram</i>	
Low-Cost Approximation Algorithms for Scheduling Independent Tasks on Hybrid Platforms.	232
<i>Louis-Claude Canon, Loris Marchal, and Frédéric Vivien</i>	

High Performance Architectures and Compilers

Runtime-Assisted Shared Cache Insertion Policies Based on Re-reference Intervals	247
<i>Vladimir Dimić, Miquel Moretó, Marc Casas, and Mateo Valero</i>	
Rewriting System for Profile-Guided Data Layout Transformations on Binaries	260
<i>Christopher Haine, Olivier Aumage, and Denis Barthou</i>	
Hardware Support for Scratchpad Memory Transactions on GPU Architectures	273
<i>Alejandro Villegas, Rafael Asenjo, Angeles Navarro, Oscar Plata, Rafael Ubal, and David Kaeli</i>	

Parallel and Distributed Data Management and Analytics

Execution of Recursive Queries in Apache Spark	289
<i>Pavlos Katsogridakis, Sofia Papagiannaki, and Polyvios Pratikakis</i>	
Replica-Aware Partitioning Design in Parallel Database Systems.	303
<i>Liming Dong, Weidong Liu, Renchuan Li, Tiejun Zhang, and Weiguo Zhao</i>	

Cluster and Cloud Computing

A Simplified Model for Simulating the Execution of a Workflow in Cloud.	319
<i>Roland Mathá, Sasko Ristov, and Radu Prodan</i>	
Dealing with Performance Unpredictability in an Asymmetric Multicore Processor Cloud	332
<i>Boris Teabe, Patrick Lavoisier Wapet, Alain Tchana, and Daniel Hagimont</i>	
Deadline-Aware Deployment for Time Critical Applications in Clouds	345
<i>Yang Hu, Junchao Wang, Huan Zhou, Paul Martin, Arie Taal, Cees de Laat, and Zhiming Zhao</i>	
More Sharing, More Benefits? A Study of Library Sharing in Container-Based Infrastructures.	358
<i>José Bravo Ferreira, Marco Cello, and Jesús Omana Iglesias</i>	
An Efficient Communication Aware Heuristic for Multiple Cloud Application Placement.	372
<i>Pedro Silva and Christian Perez</i>	

Energy-Driven Straggler Mitigation in MapReduce	385
<i>Tien-Dat Phan, Shadi Ibrahim, Amelie Chi Zhou, Guillaume Aupy, and Gabriel Antoniu</i>	

Leveraging Cloud Heterogeneity for Cost-Efficient Execution of Parallel Applications	399
<i>Eduardo Roloff, Matthias Diener, Emmanuell Diaz Carreño, Luciano Paschoal Gaspary, and Philippe O.A. Navaux</i>	

Distributed Systems and Algorithms

A Consensus-Based Fault-Tolerant Event Logger for High Performance Applications.	415
<i>Edson Tavares de Camargo, Elias P. Duarte Jr., and Fernando Pedone</i>	

Families of Graph Algorithms: SSSP Case Study	428
<i>Thejaka Amila Kanewala, Marcin Zalewski, and Andrew Lumsdaine</i>	

SEMem: Deployment of MPI-Based In-Memory Storage for Hadoop on Supercomputers	442
<i>Thanh-Chung Dao and Shigeru Chiba</i>	

Parallel and Distributed Programming, Interfaces, and Languages

Supporting the Xeon Phi Coprocessor in a Heterogeneous Programming Model	457
<i>Ana Moreton-Fernandez, Eduardo Rodriguez-Gutierrez, Arturo Gonzalez-Escribano, and Diego R. Llanos</i>	

GLT: A Unified API for Lightweight Thread Libraries	470
<i>Adrián Castelló, Sangmin Seo, Rafael Mayo, Pavan Balaji, Enrique S. Quintana-Ortí, and Antonio J. Peña</i>	

PASCAL: A Parallel Algorithmic SCALable Framework for <i>N</i> -body Problems	482
<i>Laleh Aghababaie Beni and Aparna Chandramowliswaran</i>	

GASPI/GPI In-memory Checkpointing Library	497
<i>Valeria Bartsch, Rui Machado, Dirk Merten, Mirko Rahn, and Franz-Josef Pfreundt</i>	

Multicore and Manycore Parallelism

Optimized Batched Linear Algebra for Modern Architectures	511
<i>Jack Dongarra, Sven Hammarling, Nicholas J. Higham, Samuel D. Relton, and Mawussi Zounon</i>	

New Efficient General Sparse Matrix Formats for Parallel SpMV Operations	523
<i>Jan Philipp Ecker, Rudolf Berrendorf, and Florian Mannuss</i>	
Lazy Parallel Kronecker Algebra-Operations on Heterogeneous Multicores. . .	538
<i>Wasuwee Sodsong, Robert Mittermayr, Yoojin Park, Bernd Burgstaller, and Johann Blieberger</i>	
Performance Evaluation of Computation and Communication Kernels of the Fast Multipole Method on Intel Manycore Architecture	553
<i>Mustafa Abduljabbar, Mohammed Al Farhan, Rio Yokota, and David Keyes</i>	
Efficient Non-blocking Radix Trees.	565
<i>Varun Velamuri</i>	
A Concurrency-Optimal Binary Search Tree.	580
<i>Vitaly Aksenov, Vincent Gramoli, Petr Kuznetsov, Anna Malova, and Srivatsan Ravi</i>	
Scalable Fine-Grained Metric-Based Remeshing Algorithm for Manycore/NUMA Architectures.	594
<i>Hoby Rakotoarivelo, Franck Ledoux, Franck Pommereau, and Nicolas Le-Goff</i>	
Performance Evaluation of Thread-Level Speculation in Off-the-Shelf Hardware Transactional Memories.	607
<i>Juan Salamanca, José Nelson Amaral, and Guido Araujo</i>	
Theory and Algorithms for Parallel Computation and Networking	
Addressing Volume and Latency Overheads in 1D-parallel Sparse Matrix-Vector Multiplication	625
<i>Seher Acer, Oguz Selvitopi, and Cevdet Aykanat</i>	
Improving the Network of Search Engine Services Through Application-Driven Routing	638
<i>Joe Carrión, Daniel Franco, Veronica Gil-Costa, Mauricio Marin, and Emilio Luque</i>	
Parallel Numerical Methods and Applications	
Accelerating the Tucker Decomposition with Compressed Sparse Tensors . . .	653
<i>Shaden Smith and George Karypis</i>	
Shared Memory Pipelined Parareal	669
<i>Daniel Ruprecht</i>	

Nonintrusive AMR Asynchrony for Communication Optimization. 682
*Muhammad Nufail Farooqi, Didem Unat, Tan Nguyen, Weiqun Zhang,
Ann Almgren, and John Shalf*

Accelerator Computing

Balanced CSR Sparse Matrix-Vector Product on Graphics Processors 697
Goran Flegar and Enrique S. Quintana-Ortí

To Distribute or Not to Distribute: The Question of Load Balancing
for Performance or Energy 710
*Esteban Stafford, Borja Pérez, Jose Luis Bosque, Ramón Beivide,
and Mateo Valero*

Author Index 723

Euro-Par 2017: Parallel Processing
23rd International Conference on Parallel and
Distributed Computing, Santiago de Compostela, Spain,
August 28 – September 1, 2017, Proceedings
Rivera, F.F.; Pena, T.F.; Cabaleiro, J.C. (Eds.)
2017, XXXII, 725 p. 261 illus., Softcover
ISBN: 978-3-319-64202-4