

# Contents

Asynchronous Iterative Algorithm for Computing Incomplete Factorizations on GPUs . . . . .	1
<i>Edmond Chow, Hartwig Anzt, and Jack Dongarra</i>	
Matrix Multiplication on High-Density Multi-GPU Architectures: Theoretical and Experimental Investigations . . . . .	17
<i>Peng Zhang and Yuxiang Gao</i>	
A Framework for Batched and GPU-Resident Factorization Algorithms Applied to Block Householder Transformations . . . . .	31
<i>Azzam Haidar, Tingxing Tim Dong, Stanimire Tomov, Piotr Luszczek, and Jack Dongarra</i>	
Parallel Efficient Sparse Matrix-Matrix Multiplication on Multicore Platforms . . . . .	48
<i>Md. Mostofa Ali Patwary, Nadathur Rajagopalan Satish, Narayanan Sundaram, Jongsoo Park, Michael J. Anderson, Satya Gautam Vadlamudi, Dipankar Das, Sergey G. Pudov, Vadim O. Pirogov, and Pradeep Dubey</i>	
On the Design, Development, and Analysis of Optimized Matrix-Vector Multiplication Routines for Coprocessors . . . . .	58
<i>Khairul Kabir, Azzam Haidar, Stanimire Tomov, and Jack Dongarra</i>	
Large-Scale Neo-Heterogeneous Programming and Optimization of SNP Detection on Tianhe-2 . . . . .	74
<i>Yingbo Cui, Xiangke Liao, Shaoliang Peng, Yutong Lu, Canqun Yang, Bingqiang Wang, and Chengkun Wu</i>	
ACCOLADES: A Scalable Workflow Framework for Large-Scale Simulation and Analyses of Automotive Engines . . . . .	87
<i>Shashi M. Aithal and Stefan M. Wild</i>	
Accelerating LBM and LQCD Application Kernels by In-Memory Processing . . . . .	96
<i>Paul F. Baumeister, Hans Boettiger, José R. Brunheroto, Thorsten Hater, Thilo Maurer, Andrea Nobile, and Dirk Pleiter</i>	
On Quantum Chemistry Code Adaptation for RSC PetaStream Architecture . . . . .	113
<i>Vladimir Mironov, Maria Khrenova, and Alexander Moskovsky</i>	

Dtree: Dynamic Task Scheduling at Petascale. . . . .	122
<i>Kiran Pannany, Sanchit Misra, Vasimuddin Md., Xing Liu, Edmond Chow, and Srinivas Aluru</i>	
Feasibility Study of Porting a Particle Transport Code to FPGA . . . . .	139
<i>Iakovos Panourgias, Michele Weiland, Mark Parsons, David Turland, Dave Barrett, and Wayne Gaudin</i>	
A Scalable, Linear-Time Dynamic Cutoff Algorithm for Molecular Dynamics. . . . .	155
<i>Paul Springer, Ahmed E. Ismail, and Paolo Bientinesi</i>	
BWTCP: A Parallel Method for Constructing BWT in Large Collection of Genomic Reads. . . . .	171
<i>Heng Wang, Shaoliang Peng, Yutong Lu, Chengkun Wu, Jiajun Wen, Jie Liu, and Xiaoqian Zhu</i>	
Lattice-CSC: Optimizing and Building an Efficient Supercomputer for Lattice-QCD and to Achieve First Place in Green500 . . . . .	179
<i>David Rohr, Matthias Bach, Gvozden Nešković, Volker Lindenstruth, Christopher Pinke, and Owe Philipsen</i>	
An Efficient Clique-Based Algorithm of Compute Nodes Allocation for In-memory Checkpoint System . . . . .	197
<i>Xiangke Liao, Canqun Yang, Zhe Quan, Tao Tang, and Cheng Chen</i>	
A Scalable Algorithm for Radiative Heat Transfer Using Reverse Monte Carlo Ray Tracing. . . . .	212
<i>Alan Humphrey, Todd Harman, Martin Berzins, and Phillip Smith</i>	
Optimizing Processes Mapping for Tasks with Non-uniform Data Exchange Run on Cluster with Different Interconnects . . . . .	231
<i>Victor Getmanskiy, Vladimir Chalyshev, Dmitriy Kryzhanovsky, Igor Lopatin, and Evgeny Leksikov</i>	
Dynamically Adaptable I/O Semantics for High Performance Computing . . . . .	240
<i>Michael Kuhn</i>	
Predicting Performance of Non-contiguous I/O with Machine Learning . . . . .	257
<i>Julian Kunkel, Michaela Zimmer, and Eugen Betke</i>	
A Best Practice Analysis of HDF5 and NetCDF-4 Using Lustre . . . . .	274
<i>Christopher Bartz, Konstantinos Chasapis, Michael Kuhn, Petra Nerge, and Thomas Ludwig</i>	
Striping Layout Aware Data Aggregation for High Performance I/O on a Lustre File System. . . . .	282
<i>Yuichi Tsujita, Atsushi Hori, and Yutaka Ishikawa</i>	

Hop: Elastic Consistency for Exascale Data Stores . . . . . 291  
*Latchesar Ionkov and Michael Lang*

Energy-Efficient Data Processing Through Data Sparsing with Artifacts. . . . . 307  
*Pablo Graubner, Patrick Heckmann, and Bernd Freisleben*

Updating the Energy Model for Future Exascale Systems. . . . . 323  
*Peter M. Kogge*

High-Order ADER-DG Minimizes Energy- and Time-to-Solution  
of SeisSol . . . . . 340  
*Alexander Breuer, Alexander Heinecke, Leonhard Rannabauer,  
and Michael Bader*

Modeling the Productivity of HPC Systems on a Computing Center Scale . . . 358  
*Sandra Wienke, Hristo Iliev, Dieter an Mey, and Matthias S. Müller*

Taking Advantage of Node Power Variation in Homogenous HPC Systems  
to Save Energy . . . . . 376  
*Torsten Wilde, Axel Auweter, Hayk Shoukourian, and Arndt Bode*

A Run-Time System for Power-Constrained HPC Applications. . . . . 394  
*Aniruddha Marathe, Peter E. Bailey, David K. Lowenthal,  
Barry Rountree, Martin Schulz, and Bronis R. de Supinski*

A Machine Learning Approach for a Scalable, Energy-Efficient  
Utility-Based Cache Partitioning . . . . . 409  
*Isa Ahmet Guney, Abdullah Yildiz, Ismail Ugur Bayindir,  
Kemal Cagri Serdaroglu, Utku Bayik, and Gurhan Kucuk*

A Case Study - Cost of Preemption for Urgent Computing on SuperMUC . . . 422  
*Siew Hoon Leong and Dieter Kranzlmüller*

Designing Non-blocking Personalized Collectives with Near Perfect  
Overlap for RDMA-Enabled Clusters. . . . . 434  
*Hari Subramoni, Ammar Ahmad Awan, Khaled Hamidouche,  
Dmitry Pekurovsky, Akshay Venkatesh, Sourav Chakraborty,  
Karen Tomko, and Dhabaleswar K. Panda*

Design Methodology for Optimizing Optical Interconnection Networks  
in High Performance Systems. . . . . 454  
*Sébastien Rumley, Madeleine Glick, Simon D. Hammond,  
Arun Rodrigues, and Keren Bergman*

Quantifying Communication in Graph Analytics . . . . . 472  
*Andreea Anghel, German Rodriguez, Bogdan Prisacari,  
Cyriel Minckenberg, and Gero Dittmann*

Formal Metrics for Large-Scale Parallel Performance. . . . .	488
<i>Kenneth Moreland and Ron Oldfield</i>	
Hunting Down Load Imbalance: A Moving Target . . . . .	497
<i>Christoph Pospiech</i>	
Orchestrating Docker Containers in the HPC Environment . . . . .	506
<i>Joshua Higgins, Violeta Holmes, and Colin Venters</i>	
Performance and Scaling of WRF on Three Different Parallel Supercomputers . . . . .	514
<i>Zaphiris Christidis</i>	
<b>Author Index</b> . . . . .	529



<http://www.springer.com/978-3-319-20118-4>

High Performance Computing  
30th International Conference, ISC High Performance  
2015, Frankfurt, Germany, July 12-16, 2015,  
Proceedings  
Kunkel, J.; Ludwig, Th. (Eds.)  
2015, XII, 530 p. 237 illus., Softcover  
ISBN: 978-3-319-20118-4