

Contents

Autotuning and Thread Mapping

An Analytical Model-Based Auto-tuning Framework for Locality-Aware Loop Scheduling	3
<i>Rengan Xu, Sunita Chandrasekaran, Xiaonan Tian, and Barbara Chapman</i>	
Performance, Design, and Autotuning of Batched GEMM for GPUs	21
<i>Ahmad Abdelfattah, Azzam Haidar, Stanimire Tomov, and Jack Dongarra</i>	
TCU: A Multi-Objective Hardware Thread Mapping Unit for HPC Clusters . . .	39
<i>Ravi Kumar Pujari, Thomas Wild, and Andreas Herkersdorf</i>	

Data Locality and Decomposition

Dynamic Sparse-Matrix Allocation on GPUs	61
<i>James King, Thomas Gilray, Robert M. Kirby, and Matthew Might</i>	
An Efficient Parallel Load-Balancing Framework for Orthogonal Decomposition of Geometrical Data	81
<i>Bruno R.C. Magalhães, Farhan Tauheed, Thomas Heinis, Anastasia Ailamaki, and Felix Schürmann</i>	
Parallel Community Detection Algorithm Using a Data Partitioning Strategy with Pairwise Subdomain Duplication	98
<i>Diana Palsetia, William Hendrix, Sunwoo Lee, Ankit Agrawal, Wei-keng Liao, and Alok Choudhary</i>	
TiDA: High-Level Programming Abstractions for Data Locality Management	116
<i>Didem Unat, Tan Nguyen, Weiqun Zhang, Muhammed Nufail Farooqi, Burak Bastem, George Micheliogiannakis, Ann Almgren, and John Shalf</i>	

Scalable Applications

OpenAtom: Scalable Ab-Initio Molecular Dynamics with Diverse Capabilities	139
<i>Nikhil Jain, Eric Bohm, Eric Mikida, Subhasish Mandal, Minjung Kim, Prateek Jindal, Qi Li, Sohrab Ismail-Beigi, Glenn J. Martyna, and Laxmikant V. Kale</i>	

SPRITE: A Fast Parallel SNP Detection Pipeline	159
<i>Vasudevan Rengasamy and Kamesh Madduri</i>	

Machine Learning

Predictive Modeling for Job Power Consumption in HPC Systems	181
<i>Andrea Borghesi, Andrea Bartolini, Michele Lombardi, Michela Milano, and Luca Benini</i>	
Towards Machine Learning on the Automata Processor	200
<i>Tommy Tracy II, Yao Fu, Indranil Roy, Eric Jonas, and Paul Glendenning</i>	
AutoMOMML: Automatic Multi-objective Modeling with Machine Learning	219
<i>Prasanna Balaprakash, Ananta Tiwari, Stefan M. Wild, Laura Carrington, and Paul D. Hovland</i>	

Datacenters and Cloud

Supercomputing Centers and Electricity Service Providers: A Geographically Distributed Perspective on Demand Management in Europe and the United States	243
<i>Tapasya Patki, Natalie Bates, Girish Ghatikar, Anders Clausen, Sonja Klingert, Ghaleb Abdulla, and Mehdi Sheikhalishahi</i>	
Resource Management for Running HPC Applications in Container Clouds . . .	261
<i>Stephen Herbein, Ayush Dusia, Aaron Landwehr, Sean McDaniel, Jose Monsalve, Yang Yang, Seetharami R. Seelam, and Michela Taufer</i>	

Communication Runtime

Mitigating MPI Message Matching Misery	281
<i>Mario Flajslik, James Dinan, and Keith D. Underwood</i>	
INAM ² : InfiniBand Network Analysis and Monitoring with MPI	300
<i>Hari Subramoni, Albert Mathews Augustine, Mark Arnold, Jonathan Perkins, Xiaoyi Lu, Khaled Hamidouche, and Dhableswar K. Panda</i>	
Comparing Runtime Systems with Exascale Ambitions Using the Parallel Research Kernels	321
<i>Rob F. Van der Wijngaart, Abdullah Kayi, Jeff R. Hammond, Gabriele Jost, Tom St. John, Srinivas Sridharan, Timothy G. Mattson, John Abercrombie, and Jacob Nelson</i>	

Intel Xeon Phi

High Order Seismic Simulations on the Intel Xeon Phi Processor (Knights Landing)	343
<i>Alexander Heinecke, Alexander Breuer, Michael Bader, and Pradeep Dubey</i>	
Leveraging a Cluster-Booster Architecture for Brain-Scale Simulations	363
<i>Pramod Kumbhar, Michael Hines, Aleksandr Ovcharenko, Damian A. Mallon, James King, Florentino Sainz, Felix Schürmann, and Fabien Delalondre</i>	

Manycore Architectures

Efficient and Predictable Group Communication for Manycore NoCs	383
<i>Karthik Yagna, Onkar Patil, and Frank Mueller</i>	
Distributed Job Allocation for Large-Scale Manycores	404
<i>Subramanian Ramachandran and Frank Mueller</i>	

Extreme-Scale Computations

Many-Core Acceleration of a Discrete Ordinates Transport Mini-App at Extreme Scale	429
<i>Tom Deakin, Simon McIntosh-Smith, and Wayne Gaudin</i>	
Efficiency of High Order Spectral Element Methods on Petascale Architectures	449
<i>Maxwell Hutchinson, Alexander Heinecke, Hans Pabst, Greg Henry, Matteo Parsani, and David Keyes</i>	

Resilience

Scalability of Partial Differential Equations Preconditioner Resilient to Soft and Hard Faults	469
<i>Karla Morris, Francesco Rizzi, Khachik Sargsyan, Kathryn Dahlgren, Paul Mycek, Cosmin Safta, Olivier Le Maître, Omar Knio, and Bert Debusschere</i>	
Multi-versioning Performance Opportunities in BGAS System for Resilience	486
<i>Nan Dun, Dirk Pleiter, Aiman Fang, Nicolas Vandenberg, and Andrew A. Chien</i>	

Author Index	505
-------------------------------	-----

High Performance Computing

31st International Conference, ISC High Performance

2016, Frankfurt, Germany, June 19-23, 2016,

Proceedings

Kunkel, J.M.; Balaji, P.; Dongarra, J. (Eds.)

2016, XV, 506 p. 227 illus., Softcover

ISBN: 978-3-319-41320-4