

Contents

Invited Papers

Concurrent Systems: Hybrid Object Implementations and Abortable Objects	3
<i>Michel Raynal</i>	

Runtime-Aware Architectures	16
<i>Marc Casas, Miquel Moreto, Lluc Alvarez, Emilio Castillo, Dimitrios Chasapis, Timothy Hayes, Luc Jaulmes, Oscar Palomar, Osman Unsal, Adrian Cristal, Eduard Ayguade, Jesus Labarta, and Mateo Valero</i>	

Support Tools and Environments

MPI Thread-Level Checking for MPI+OpenMP Applications	31
<i>Emmanuelle Saillard, Patrick Carribault, and Denis Barthou</i>	

Event-Action Mappings for Parallel Tools Infrastructures	43
<i>Tobias Hilbrich, Martin Schulz, Holger Brunst, Joachim Protze, Bronis R. de Supinski, and Matthias S. Müller</i>	

Performance Modeling, Prediction and Evaluation

Low-Overhead Detection of Memory Access Patterns and Their Time Evolution	57
<i>Harald Servat, Germán Llort, Juan González, Judit Giménez, and Jesús Labarta</i>	

Automatic On-Line Detection of MPI Application Structure with Event Flow Graphs	70
<i>Xavier Aguilar, Karl Furlinger, and Erwin Laure</i>	

Online Automated Reliability Classification of Queueing Models for Streaming Processing Using Support Vector Machines	82
<i>Jonathan C. Beard, Cooper Epstein, and Roger D. Chamberlain</i>	

Scheduling and Load Balancing

A Duplicate-Free State-Space Model for Optimal Task Scheduling	97
<i>Michael Orr and Oliver Sinnen</i>	

On the Heterogeneity Bias of Cost Matrices When Assessing Scheduling Algorithms.	109
<i>Louis-Claude Canon and Laurent Philippe</i>	
Hardware Round-Robin Scheduler for Single-ISA Asymmetric Multi-core . . .	122
<i>Nikola Markovic, Daniel Nemirovsky, Veljko Milutinovic, Osman Unsal, Mateo Valero, and Adrian Cristal</i>	
Moody Scheduling for Speculative Parallelization	135
<i>Alvaro Estebanez, Diego R. Llanos, David Orden, and Belen Palop</i>	
Allocating Jobs with Periodic Demand Variations	147
<i>Olivier Beaumont, Ikbel Belaid, Lionel Eyraud-Dubois, and Juan-Angel Lorenzo-del-Castillo</i>	
A Multi-level Hypergraph Partitioning Algorithm Using Rough Set Clustering.	159
<i>Foad Lotfifar and Matthew Johnson</i>	
Non-preemptive Throughput Maximization for Speed-Scaling with Power-Down.	171
<i>Eric Angel, Evripidis Bampis, Vincent Chau, and Nguyen Kim Thang</i>	
Scheduling Tasks from Selfish Multi-tasks Agents	183
<i>Johanne Cohen and Fanny Pascual</i>	
Locality and Balance for Communication-Aware Thread Mapping in Multicore Systems.	196
<i>Matthias Diener, Eduardo H.M. Cruz, Marco A.Z. Alves, Mohammad S. Alhakeem, Philippe O.A. Navaux, and Hans-Ulrich Hei</i>	
Priority Queues Are Not Good Concurrent Priority Schedulers	209
<i>Andrew Lenharth, Donald Nguyen, and Keshav Pingali</i>	
Load Balancing Prioritized Tasks via Work-Stealing	222
<i>Shams Imam and Vivek Sarkar</i>	
Architecture and Compilers	
Optimizing Task Parallelism with Library-Semantics-Aware Compilation. . . .	237
<i>Peter Thoman, Stefan Moosbrugger, and Thomas Fahringer</i>	
Data Layout Optimization for Portable Performance	250
<i>Kamal Sharma, Ian Karlin, Jeff Keasler, James R. McGraw, and Vivek Sarkar</i>	
Automatic Data Layout Optimizations for GPUs.	263
<i>Klaus Kofler, Biagio Cosenza, and Thomas Fahringer</i>	

Parallel and Distributed Data Management

Performance Impacts with Reliable Parallel File Systems at Exascale Level. . .	277
<i>Ramon Nou, Alberto Miranda, and Toni Cortes</i>	
Rapid Tomographic Image Reconstruction via Large-Scale Parallelization . . .	289
<i>Tekin Bicer, Doga Gursoy, Rajkumar Kettimuthu, Francesco De Carlo, Gagan Agrawal, and Ian T. Foster</i>	

Grid, Cluster and Cloud Computing

Software Consolidation as an Efficient Energy and Cost Saving Solution for a SaaS/PaaS Cloud Model.	305
<i>Alain Tchana, Noel De Palma, Ibrahim Safieddine, Daniel Hagimont, Bruno Diot, and Nicolas Vuillerme</i>	
VMPlaceS: A Generic Tool to Investigate and Compare VM Placement Algorithms.	317
<i>Adrien Lebre, Jonathan Pastor, and Mario Südholt</i>	

Distributed Systems and Algorithms

A Connectivity Model for Agreement in Dynamic Systems	333
<i>Carlos Gómez-Calzado, Arnaud Casteigts, Alberto Lafuente, and Mikel Larrea</i>	
DFEP: Distributed Funding-Based Edge Partitioning	346
<i>Alessio Guerrieri and Alberto Montresor</i>	

Parallel and Distributed Programming, Interfaces and Languages

PR-STM: Priority Rule Based Software Transactions for the GPU.	361
<i>Qi Shen, Craig Sharp, William Blewitt, Gary Ushaw, and Graham Morgan</i>	
Leveraging MPI-3 Shared-Memory Extensions for Efficient PGAS Runtime Systems	373
<i>Huan Zhou, Kamran Idrees, and José Gracia</i>	

Multi- and Many-core Programming

A Practical Transactional Memory Interface	387
<i>Shahar Timnat, Maurice Herlihy, and Erez Petrank</i>	
A Multicore Parallelization of Continuous Skyline Queries on Data Streams	402
<i>Tiziano De Matteis, Salvatore Di Girolamo, and Gabriele Mencagli</i>	

A Fast and Scalable Graph Coloring Algorithm for Multi-core and Many-core Architectures	414
<i>Georgios Rokos, Gerard Gorman, and Paul H.J. Kelly</i>	
A Composable Deadlock-Free Approach to Object-Based Isolation	426
<i>Shams Imam, Jisheng Zhao, and Vivek Sarkar</i>	
Scalable Data-Driven PageRank: Algorithms, System Issues, and Lessons Learned	438
<i>Joyce Jiyoung Whang, Andrew Lenharth, Inderjit S. Dhillon, and Keshav Pingali</i>	
How Many Threads will be too Many? On the Scalability of OpenMP Implementations	451
<i>Christian Iwainsky, Sergei Shudler, Alexandru Calotoiu, Alexandre Strube, Michael Knobloch, Christian Bischof, and Felix Wolf</i>	
Theory and Algorithms for Parallel Computation	
Efficient Nested Dissection for Multicore Architectures	467
<i>Dominique LaSalle and George Karypis</i>	
Scheduling Trees of Malleable Tasks for Sparse Linear Algebra	479
<i>Abdou Guermouche, Loris Marchal, Bertrand Simon, and Frédéric Vivien</i>	
Elastic Tasks: Unifying Task Parallelism and SPMD Parallelism with an Adaptive Runtime	491
<i>Alina Sbîrlea, Kunal Agrawal, and Vivek Sarkar</i>	
Numerical Methods and Applications	
Semi-discrete Matrix-Free Formulation of 3D Elastic Full Waveform Inversion Modeling	507
<i>Stephen Moore, Devi Sudheer Chunduri, Sergiy Zhuk, Tigran Tchraikian, Ewout van den Berg, Albert Akhriev, Alberto Costa Nogueira Jr., Andrew Rawlinson, and Lior Horeish</i>	
10,000 Performance Models per Minute – Scalability of the UG4 Simulation Framework.	519
<i>Andreas Vogel, Alexandru Calotoiu, Alexandre Strube, Sebastian Reiter, Arne Nägel, Felix Wolf, and Gabriel Wittum</i>	
Exploiting Task-Based Parallelism in Bayesian Uncertainty Quantification . . .	532
<i>Panagiotis E. Hadjidoukas, Panagiotis Angelikopoulos, Lina Kulakova, Costas Papadimitriou, and Petros Koumoutsakos</i>	

Parallelization of an Advection-Diffusion Problem Arising in Edge Plasma Physics Using Hybrid MPI/OpenMP Programming	545
<i>Matthieu Kuhn, Guillaume Latu, Nicolas Crouseilles, and Stéphane Genaud</i>	
Behavioral Non-portability in Scientific Numeric Computing	558
<i>Yijia Gu, Thomas Wahl, Mahsa Bayati, and Miriam Leeser</i>	
Accelerator Computing	
Fast Parallel Suffix Array on the GPU.	573
<i>Leyuan Wang, Sean Baxter, and John D. Owens</i>	
Effective Barrier Synchronization on Intel Xeon Phi Coprocessor	588
<i>Andrey Rodchenko, Andy Nisbet, Antoniu Pop, and Mikel Luján</i>	
High Performance Multi-GPU SpMV for Multi-component PDE-Based Applications	601
<i>Ahmad Abdelfattah, Hatem Ltaief, and David Keyes</i>	
Accelerating Lattice Boltzmann Applications with OpenACC	613
<i>Enrico Calore, Jiri Kraus, Sebastiano Fabio Schifano, and Raffaele Tripiccone</i>	
High-Performance and Scalable Design of MPI-3 RMA on Xeon Phi Clusters	625
<i>Mingzhe Li, Khaled Hamidouche, Xiaoyi Lu, Jian Lin, and Dhabaleswar K. (DK) Panda</i>	
Improving Performance of Convolutional Neural Networks by Separable Filters on GPU	638
<i>Hao-Ping Kang and Che-Rung Lee</i>	
Iterative Sparse Triangular Solves for Preconditioning	650
<i>Hartwig Anzt, Edmond Chow, and Jack Dongarra</i>	
Targeting the Parallella	662
<i>Spiros N. Agathos, Alexandros Papadogiannakis, and Vassilios V. Dimakopoulos</i>	
Systematic Fusion of CUDA Kernels for Iterative Sparse Linear System Solvers	675
<i>José I. Aliaga, Joaquín Pérez, and Enrique S. Quintana-Ortí</i>	
Efficient Execution of Multiple CUDA Applications Using Transparent Suspend, Resume and Migration.	687
<i>Taichiro Suzuki, Akira Nukada, and Satoshi Matsuoka</i>	
Author Index	701