

Table of Contents

Invited Talk

Selfish Distributed Optimization	1
<i>Burkhard Monien and Christian Scheideler</i>	

Topic 1: Support Tools and Environments

Introduction.....	3
<i>Omer Rana, Marios Dikaiakos, Daniel S. Katz, and Christine Morin</i>	
Tulip: A Visualization Framework for User-Guided Parallelization	4
<i>Yi Wen Wong, Tomasz Dubrownik, Wai Teng Tang, Wen Jun Tan, Rubing Duan, Rick Siow Mong Goh, Shyh-hao Kuo, Stephen John Turner, and Weng-Fai Wong</i>	
Enabling Cloud Interoperability with COMPSs	16
<i>Fabrizio Marozzo, Francesc Lordan, Roger Rafanell, Daniele Lezzi, Domenico Talia, and Rosa M. Badia</i>	
Pattern-Independent Detection of Manual Collectives in MPI Programs	28
<i>Alexandru Calotoiu, Christian Siebert, and Felix Wolf</i>	
A Type-Based Approach to Separating Protocol from Application Logic: A Case Study in Hybrid Computer Programming	40
<i>Geoffrey C. Hulette, Matthew J. Sottile, and Allen D. Malony</i>	

Topic 2: Performance Prediction and Evaluation

Introduction.....	52
<i>Allen D. Malony, Helen Karatza, William Knottenbelt, and Sally McKee</i>	
Energy Consumption Modeling for Hybrid Computing	54
<i>Ami Marowka</i>	
HPC File Systems in Wide Area Networks: Understanding the Performance of Lustre over WAN	65
<i>Alvaro Aguilera, Michael Kluge, Thomas William, and Wolfgang E. Nagel</i>	

Understanding I/O Performance Using I/O Skeletal Applications	77
<i>Jeremy Logan, Scott Klasky, Hasan Abbasi, Qing Liu, George Ostrouchov, Manish Parashar, Norbert Podhorszki, Yuan Tian, and Matthew Wolf</i>	
ASK: Adaptive Sampling Kit for Performance Characterization	89
<i>Pablo de Oliveira Castro, Eric Petit, Jean Christophe Beyler, and William Jalby</i>	
CRAW/P: A Workload Partition Method for the Efficient Parallel Simulation of Manycores	102
<i>Shuai Jiao, Paolo Ienne, Xiaochun Ye, Da Wang, Dongrui Fan, and Ninghui Sun</i>	
Topic 3: Scheduling and Load Balancing	
Introduction	115
<i>Denis Trystram, Ioannis Milis, Zhihui Du, and Uwe Schwiegelshohn</i>	
Job Scheduling Using Successive Linear Programming Approximations of a Sparse Model	116
<i>Stephane Chretien, Jean-Marc Nicod, Laurent Philippe, Veronika Rehn-Sonigo, and Lamiel Toch</i>	
Speed Scaling on Parallel Processors with Migration	128
<i>Eric Angel, Eiripidis Bampis, Fadi Kacem, and Dimitrios Letsios</i>	
Dynamic Distributed Scheduling Algorithm for State Space Search	141
<i>Ankur Narung, Abhinav Srivastava, Ramnik Jain, and R.K. Shyamasundar</i>	
Using Load Information in Work-Stealing on Distributed Systems with Non-uniform Communication Latencies	155
<i>Vladimir Janjic and Kevin Hammond</i>	
Energy Efficient Frequency Scaling and Scheduling for Malleable Tasks	167
<i>Peter Sanders and Jochen Speck</i>	
Scheduling MapReduce Jobs in HPC Clusters	179
<i>Marcelo Veiga Neves, Tiago Ferreto, and César De Rose</i>	
A Job Scheduling Approach for Multi-core Clusters Based on Virtual Malleability	191
<i>Gladys Utrera, Siham Tabik, Julita Corbalan, and Jesús Labarta</i>	

Topic 4: High-Performance Architecture and Compilers

Introduction.....	204
<i>Alex Veidenbaum, Nectarios Koziris, Toshinori Sato, and Avi Mendelson</i>	
Dynamic Last-Level Cache Allocation to Reduce Area and Power Overhead in Directory Coherence Protocols	206
<i>Mario Lodde, Jose Flich, and Manuel E. Acacio</i>	
A Practical Approach to DOACROSS Parallelization	219
<i>Priya Unnikrishnan, Jun Shirako, Kit Barton, Sanjay Chatterjee, Raul Silvera, and Vivek Sarkar</i>	
Exploiting Semantics of Virtual Memory to Improve the Efficiency of the On-Chip Memory System.....	232
<i>Bin Li, Zhen Fang, Li Zhao, Xiaowei Jiang, Lin Li, Andrew Herdrich, Ravishankar Iyer, and Srihari Makineni</i>	
From Serial Loops to Parallel Execution on Distributed Systems	246
<i>George Bosilca, Aurelien Bouteiller, Anthony Danalis, Thomas Herault, and Jack Dongarra</i>	

Topic 5: Parallel and Distributed Data Management

Introduction.....	258
<i>Domenico Talia, Alex Delis, Haimonti Dutta, and Arkady Zaslavsky</i>	
DS-Means: Distributed Data Stream Clustering	260
<i>Alessio Guerrieri and Alberto Montresor</i>	
3D Inverted Index with Cache Sharing for Web Search Engines	272
<i>Esteban Feuerstein, Veronica Gil-Costa, Mauricio Marin, Gabriel Tolosa, and Ricardo Baeza-Yates</i>	
Quality-of-Service for Consistency of Data Geo-replication in Cloud Computing.....	285
<i>Sérgio Esteves, João Silva, and Luís Veiga</i>	
A Fault-Tolerant Cache Service for Web Search Engines: RADIC Evaluation	298
<i>Carlos Gómez-Pantoja, Dolores Rexachs, Mauricio Marin, and Emilio Luque</i>	

Topic 6: Grid, Cluster and Cloud Computing

Introduction..... 311
*Erik Elmroth, Paraskevi Fragopoulou, Artur Andrzejak,
Ivona Brandic, Karim Djemame, and Paolo Romano*

Scalable Reed-Solomon-Based Reliable Local Storage for HPC
Applications on IaaS Clouds..... 313
*Leonardo Bautista Gomez, Bogdan Nicolae, Naoya Maruyama,
Franck Cappello, and Satoshi Matsuoka*

Caching VM Instances for Fast VM Provisioning: A Comparative
Evaluation 325
Pradipta De, Manish Gupta, Manoj Soni, and Aditya Thatte

Improving Scheduling Performance Using a Q-Learning-Based Leasing
Policy for Clouds 337
Alexander Fölling and Matthias Hofmann

Impact of Variable Priced Cloud Resources on Scientific Workflow
Scheduling 350
Simon Ostermann and Radu Prodan

Topic 7: Peer to Peer Computing

Introduction..... 363
*Alberto Montresor, Evaggelia Pitoura, Anwitaman Datta, and
Spyros Voulgaris*

ID-Replication for Structured Peer-to-Peer Systems 364
Tallat M. Shafaat, Bilal Ahmad, and Seif Haridi

Changing the Unchoking Policy for an Enhanced Bittorrent 377
Vaggelis Atlidakis, Mema Roussopoulos, and Alex Delis

Peer-to-Peer Multi-class Boosting 389
*István Hegedűs, Róbert Busa-Fekete, Róbert Ormándi,
Márk Jelasity, and Balázs Kégl*

Topic 8: Distributed Systems and Algorithms

Introduction..... 401
*Andrzej Goscinski, Marios Mavronicolas, Weisong Shi, and
Teo Yong Meng*

Towards Load Balanced Distributed Transactional Memory 403
Gokarna Sharma and Costas Busch

CUDA-For-Clusters: A System for Efficient Execution of CUDA Kernels on Multi-core Clusters	415
<i>Raghu Prabhakar, R. Govindarajan, and Matthew J. Thazhuthaveetil</i>	
From a Store-Collect Object and Ω to Efficient Asynchronous Consensus	427
<i>Michel Raynal and Julien Stainer</i>	
An Investigation into the Performance of Reduction Algorithms under Load Imbalance	439
<i>Petar Marendić, Jan Lemeire, Tom Haber, Dean Vučinić, and Peter Schelkens</i>	
Achieving Reliability in Master-Worker Computing via Evolutionary Dynamics	451
<i>Eugenia Christoforou, Antonio Fernández Anta, Chryssis Georgiou, Miguel A. Mosteiro, and Angel (Anxo) Sánchez</i>	
Topic 9: Parallel and Distributed Programming	
Introduction	464
<i>Sergei Gorlatch, Rizos Sakellariou, Marco Danelutto, and Thilo Kielmann</i>	
Dynamic Thread Mapping Based on Machine Learning for Transactional Memory Applications	465
<i>Márcio Castro, Luís Fabrício Wanderley Góes, Luiz Gustavo Fernandes, and Jean-François Méhaut</i>	
A Checkpoint-on-Failure Protocol for Algorithm-Based Recovery in Standard MPI	477
<i>Wesley Bland, Peng Du, Aurelien Bouteiller, Thomas Herault, George Bosilca, and Jack Dongarra</i>	
Hierarchical Partitioning Algorithm for Scientific Computing on Highly Heterogeneous CPU + GPU Clusters	489
<i>David Clarke, Aleksandar Ilic, Alexey Lastovetsky, and Leonel Sousa</i>	
Encapsulated Synchronization and Load-Balance in Heterogeneous Programming	502
<i>Yuri Torres, Arturo Gonzalez-Escribano, and Diego Llanos</i>	
Transactional Access to Shared Memory in StarSs, a Task Based Programming Model	514
<i>Rahul Kumar Gayatri, Rosa M. Badia, Eduard Ayguade, Mikel Luján, and Ian Watson</i>	
On-the-Fly Task Execution for Speeding Up Pipelined MapReduce	526
<i>Diana Moise, Gabriel Antoniu, and Luc Bougé</i>	

Assessing the Performance and Scalability of a Novel Multilevel
K-Nomial Allgather on CORE-Direct Systems 538
*Joshua S. Ladd, Manjunath Gorentla Venkata,
Richard Graham, and Pavel Shamis*

Topic 10: Parallel Numerical Algorithms

Introduction..... 550
Iain Duff, Efstratios Gallopoulos, Daniela di Serafino, and Bora Ucar

Avoiding Communication through a Multilevel LU Factorization 551
Simplice Donfack, Laura Grigori, and Amal Khabou

Locality Improvement of Data-Parallel Adams–Bashforth Methods
through Block-Based Pipelining of Time Steps 563
Matthias Korch

Parallel SOR for Solving the Convection Diffusion Equation Using
GPUs with CUDA 575
*Yiannis Cotronis, Elias Konstantinidis, Maria A. Louka, and
Nikolaos M. Missirlis*

Topic 11: Multicore and Manycore Programming

Introduction..... 587
*Eduard Ayguade, Dionisios Pnevmatikatos, Rudolf Eigenmann,
Mikel Luján, and Sabri Pllana*

Efficient Support for In-Place Metadata in Transactional Memory 589
Ricardo J. Dias, Tiago M. Vale, and João M. Lourenço

Folding of Tagged Single Assignment Values for Memory-Efficient
Parallelism 601
Dragoş Sbîrlea, Kathleen Knobe, and Vivek Sarkar

High-Level Support for Pipeline Parallelism on Many-Core
Architectures 614
*Siegfried Benkner, Enès Bajrovic, Erich Marth, Martin Sandrieser,
Raymond Namyst, and Samuel Thibault*

Node.Scala: Implicit Parallel Programming for High-Performance Web
Services..... 626
*Daniele Bonetta, Danilo Ansaloni, Achille Peternier,
Cesare Pautasso, and Walter Binder*

Task-Parallel Programming on NUMA Architectures 638
Christian Terboven, Dirk Schmid, Tim Cramer, and Dieter an Mey

Speeding Up OpenMP Tasking	650
<i>Spiros N. Agathos, Nikolaos D. Kallimanis, and Vassilios V. Dimakopoulos</i>	

An Efficient Unbounded Lock-Free Queue for Multi-core Systems	662
<i>Marco Aldinucci, Marco Danehutto, Peter Kilpatrick, Massimiliano Meneghin, and Massimo Torquati</i>	

Topic 12: Theory and Algorithms for Parallel Computation

Introduction.....	674
<i>Geppino Pucci, Christos Zaroliagis, Kieran T. Herley, and Henning Meyerhenke</i>	

A Lower Bound Technique for Communication on BSP with Application to the FFT.....	676
<i>Gianfranco Bilardi, Michele Scquizzato, and Francesco Silvestri</i>	

A Fast Parallel Algorithm for Minimum-Cost Small Integral Flows	688
<i>Andrzej Lingas and Mia Persson</i>	

Topic 13: High Performance Network and Communication

Introduction.....	700
<i>Chris Develder, Emmanouel Varvarigos, Admela Jukan, and Dimitra Simeonidou</i>	

Topology Configuration in Hybrid EPS/OCS Interconnects.....	701
<i>Konstantinos Christodoulopoulos, Marco Ruffini, Donal O'Mahony, and Kostas Katrinis</i>	

Towards an Efficient Fat-Tree Like Topology	716
<i>D. Bermúdez Garzón, C. Gómez, M.E. Gómez, P. López, and J. Duato</i>	

An Adaptive, Scalable, and Portable Technique for Speeding Up MPI-Based Applications	729
<i>Rosa Filgueira, Malcolm Atkinson, Alberto Nuñez, and Javier Fernández</i>	

Cost-Effective Contention Avoidance in a CMP with Shared Memory Controllers	741
<i>Samuel Rodrigo, Frank Olaf Sem-Jacobsen, Hervé Tutenguem, Tor Skeie, and Davide Bertozzi</i>	

Topic 14: Mobile and Ubiquitous Computing

Introduction 753
Paolo Santi, Sotiris Nikolettseas, Cecilia Mascolo, and Thiemo Voigt

Watershed-Based Clustering for Energy Efficient Data Gathering in
Wireless Sensor Networks with Mobile Collector 754
*Charalampos Konstantopoulos, Basilis Mamalis,
Grammati Pantziou, and Vasileios Thanasias*

Distribution of Liveness Property Connectivity Interval in Selected
Mobility Models of Wireless Ad Hoc Networks 767
Jerzy Brzeziński, Michał Kalewski, Marcin Kosiba, and Marek Libuda

Topic 15: High Performance and Scientific
Applications

Introduction 779
Thomas Ludwig, Costas Bekas, Alice Koniges, and Kengo Nakajima

Memory-Access Optimization of Parallel Molecular Dynamics
Simulation via Dynamic Data Reordering 781
*Manaschai Kunaseth, Ken-ichi Nomura, Hikmet Dursun,
Rajiv K. Kalia, Aiichiro Nakano, and Priya Vashishta*

On Analyzing Quality of Data Influences on Performance of Finite
Elements Driven Computational Simulations 793
*Michael Reiter, Hong-Linh Truong, Schahram Dustdar,
Dimka Karastoyanova, Robert Krause, Frank Leymann, and
Dieter Pahr*

Performance Evaluation and Optimization of Nested High Resolution
Weather Simulations 805
*Preeti Malakar, Vaibhav Saxena, Thomas George, Rashmi Mittal,
Sameer Kumar, Abdul Ghani Naim, and Saiful Azmi bin Hj Husain*

Optimized Hybrid Parallel Lattice Boltzmann Fluid Flow Simulations
on Complex Geometries 818
*Jonas Fietz, Mathias J. Krause, Christian Schulz,
Peter Sanders, and Vincent Heuveline*

Topology-Aware Mappings for Large-Scale Eigenvalue Problems 830
*Hasan Metin Aktulga, Chao Yang, Esmond G. Ng,
Pieter Maris, and James P. Vary*

Fast and Effective Lossy Compression Algorithms for Scientific
Datasets 843
Jeremy Iverson, Chandrika Kamath, and George Karypis

Topic 16: GPU and Accelerators Computing

Introduction	857
<i>Alex Ramirez, Dimitrios S. Nikolopoulos, David Kaeli, and Satoshi Matsuoka</i>	
OpenACC — First Experiences with Real-World Applications	859
<i>Sandra Wienke, Paul Springer, Christian Terboven, and Dieter an Mey</i>	
accULL: An OpenACC Implementation with CUDA and OpenCL Support	871
<i>Ruymán Reyes, Iván López-Rodríguez, Juan J. Fumero, and Francisco de Sande</i>	
Understanding the Performance of Concurrent Data Structures on Graphics Processors	883
<i>Daniel Cederman, Bapi Chatterjee, and Philippas Tsigas</i>	
A New Programming Paradigm for GPGPU	895
<i>Julio Toss and Thierry Gautier</i>	
GPU-Accelerated Asynchronous Error Correction for Mixed Precision Iterative Refinement	908
<i>Hartwig Anzt, Piotr Luszczek, Jack Dongarra, and Vincent Heuveline</i>	
GPURoofline: A Model for Guiding Performance Optimizations on GPUs	920
<i>Haipeng Jia, Yunquan Zhang, Guoping Long, Jianliang Xu, Shengen Yan, and Yan Li</i>	
Building a Collision for 75-Round Reduced SHA-1 Using GPU Clusters	933
<i>Andrew V. Adinetz and Evgeny A. Grechnikov</i>	
GPU-Vote: A Framework for Accelerating Voting Algorithms on GPU	945
<i>Gert-Jan van den Braak, Cedric Nugteren, Bart Mesman, and Henk Corporaal</i>	
Author Index	957