Jesus Carretero · Javier Garcia-Blas
Ryan K.L. Ko · Peter Mueller
Koji Nakano (Eds.)

Algorithms and Architectures for Parallel Processing

16th International Conference, ICA3PP 2016
Granada, Spain, December 14–16, 2016
Proceedings

Springer

Contents

Parallel and Distributed Architectures

Intelligent SPARQL Endpoints: Optimizing Execution Performance by Automatic Query Relaxation and Queue Scheduling 3
Ana I. Torre-Bastida, Esther Villar-Rodriguez, Miren Nekane Bilbao, and Javier Del Ser

Hardware-Based Sequential Consistency Violation Detection Made Simpler . . . 18
Mohammad Majharul Islam, Riad Akram, and Abdullah Muzahid

Optimized Mapping Spiking Neural Networks onto Network-on-Chip 38
Yu Ji, Youhui Zhang, He Liu, and Weimin Zheng

Software Systems and Programming

A Portable Lock-Free Bounded Queue 55
Peter Pirkelbauer, Reed Milewicz, and Juan Felipe Gonzalez

A C++ Generic Parallel Pattern Interface for Stream Processing 74
David del Rio Astorga, Manuel F. Dolz, Luis Miguel Sanchez, Javier García Blas, and J. Daniel García

Creating Distributed Execution Plans with BobolangNG 88
David Bednárek, Martin Kruliš, Jakub Yaghob, and Filip Zavoral

Deciding the Deadlock and Livelock in a Petri Net with a Target Marking Based on Its Basic Unfolding 98
Guanjun Liu, Kun Zhang, and Changjun Jiang

A New Scalable Approach for Distributed Metadata in HPC 106
Cristina Rodríguez-Quintana, Antonio F. Díaz, Julio Ortega, Raúl H. Palacios, and Andrés Ortiz

Enabling Android-Based Devices to High-End GPGPUs 118
Raffaele Montella, Carmine Ferraro, Sokol Kosta, Valentina Pelliccia, and Giulio Giunta

Distributed and Network-Based Computing

3-Additive Approximation Algorithm for Multicast Time in 2D Torus Networks 129
Hovhaness A. Harutyunyan and Meghrig Terzian

Online Resource Coalition Reorganization for Efficient Scheduling on the Intercloud 143
Adrian Spataru, Teodora Selea, and Marc Frincu

Graphein: A Novel Optical High-Radix Switch Architecture for 3D Integration 162
Jie Jian, Mingche Lai, Liquan Xiao, and Weixia Xu

Improving the Performance of Volunteer Computing with Data Volunteers: A Case Study with the ATLAS@home Project 178
Saúl Alonso-Monsalve, Félix García-Carballeira, and Alejandro Calderón

Microcities: A Platform Based on Microclouds for Neighborhood Services . . . 192
Ismael Cuadrado-Cordero, Felix Cuadrado, Chris Phillips, Anne-Cécile Orgerie, and Christine Morin

Impact of Shutdown Techniques for Energy-Efficient Cloud Data Centers . . . 203
Issam Raïs, Anne-Cécile Orgerie, and Martin Quinson

Processing Partially Ordered Requests in Distributed Stream Processing Systems 211
Rijun Cai, Weigang Wu, Ning Huang, and Lihui Wu

Implement and Optimization of Indoor Positioning System Based on Wi-Fi Signal 220
Chongsheng Yu, Xin Li, Lei Dou, Jianwei Li, Yu Zhang, Jian Qin, Yuqing Sun, and Zhiyue Cao

Big Data and Its Applications

Optimizing Inter-server Communications by Exploiting Overlapping Communities in Online Social Networks 231
Jingya Zhou, Jianxi Fan, Baolei Cheng, and Juncheng Jia

Road Segment Information Based Named Data Networking for Vehicular Environments 245
Junlan Xiao, Jian Deng, Hui Cao, and Weigang Wu

Energy-Aware Query Processing on a Parallel Database Cluster Node 260
Amine Roukh, Ladjel Bellatreche, Nikos Tziritas, and Carlos Ordonez

Current Flow Betweenness Centrality with Apache Spark 270
Massimiliano Bertolucci, Alessandro Lulli, and Laura Ricci

Parallel and Distributed Algorithms

Light Loss-Less Data Compression, with GPU Implementation 281
Shunji Funasaka, Koji Nakano, and Yasuaki Ito

Deterministic Construction of Regular Geometric Graphs with Short Average Distance and Limited Edge Length 295
Satoshi Fujita, Koji Nakano, Michihiro Koibuchi, and Ikki Fujiwara

A GPU-Based Backtracking Algorithm for Permutation Combinatorial Problems 310
Tiago Carneiro Pessoa, Jan Gmys, Nouredine Melab, Francisco Heron de Carvalho Junior, and Daniel Tuyttens

Buffer Minimization for Rate-Optimal Scheduling of Synchronous Dataflow Graphs on Multicore Systems 325
Mingze Ma and Rizos Sakellariou

Implementing Snapshot Objects on Top of Crash-Prone Asynchronous Message-Passing Systems 341
Carole Delporte-Gallet, Hugues Fauconnier, Sergio Rajsbaum, and Michel Raynal

Scaling DBSCAN-like Algorithms for Event Detection Systems in Twitter . . . 356
Joan Capdevila, Gonzalo Pericacho, Jordi Torres, and Jesús Cerquides

Towards Parallel CFD Computation for the ADAPT Framework 374
Imad Kissami, Christophe Cérin, Fayssal Benkhaldoun, and Gilles Scarella

Feedback Control Optimization for Performance and Energy Efficiency on CPU-GPU Heterogeneous Systems 388
Feng-Sheng Lin, Po-Ting Liu, Ming-Hua Li, and Pao-Ann Hsiung

The Impact of Panel Factorization on the Gauss-Huard Algorithm for the Solution of Linear Systems on Modern Architectures 405
Sandra Catalán, Pablo Ezzatti, Enrique S. Quintana-Ortí, and Alfredo Remón

Leveraging the Performance of LBM-HPC for Large Sizes on GPUs Using Ghost Cells 417
Pedro Valero-Lara

Improving Hash Distributed A* for Shared Memory Architectures Using Abstraction 431
Victoria Sanz, Armando De Giusti, and Marcelo Naiouf

On a Parallel Algorithm for the Determination of Multiple Optimal Solutions for the LCSS Problem 440
Bchira Ben Mabrouk, Hamadi Hasni, and Zaher Mahjoub

Locality of Computation for Stencil Optimization 449
Lufeng Yuan, Junhong Liu, Yulong Luo, and Guangming Tan

GPU Computing to Speed-Up the Resolution of Microrheology Models. 457
Gloria Ortega, Antonio Puertas, Fco Javier de Las Nieves, and Ester Martin-Garzón

Applications of Parallel and Distributed Computing

Methodological Approach to Data-Centric Cloudification of Scientific Iterative Workflows. 469
Silvina Caíno-Lores, Andrei Lapin, Peter Kropf, and Jesús Carretero

Efficient Parallel Algorithm for Optimal DAG Structure Search on Parallel Computer with Torus Network 483
Hirokazu Honda, Yoshinori Tamada, and Reiji Suda

Bin Recycling Strategy for an Accuracy-Aware Implementation of Two-Point Angular Correlation Function on GPU 503
Miguel Cárdenas-Montes, Juan José Rodríguez-Vázquez, Miguel A. Vega-Rodríguez, Ignacio Sevilla Noarbe, and Antonio Gómez-Iglesias

An Efficient Implementation of LZW Compression in the FPGA 512
Xin Zhou, Yasuaki Ito, and Koji Nakano

Shared Memory Tile-Based vs Hybrid Memory GOP-Based Parallel Algorithms for HEVC Encoder. 521
Héctor Migallón, Otoniel López-Granado, Vicente Galiano, Pablo Piñol, and Manuel P. Malumbres

GPU-Based Heterogeneous Coding Architecture for HEVC 529
Gabriel Cebrián-Márquez, Héctor Migallón, José Luis Martínez, Otoniel López-Granado, Pablo Piñol, and Pedro Cuenca

Optimizing GPU Code for CPU Execution Using OpenCL and Vectorization: A Case Study on Image Coding 537
Pedro M.M. Pereira, Patricio Domingues, Nuno M.M. Rodrigues, Gabriel Falcao, and Sergio M.M. de Faria

Improving the Performance of Cardiac Simulations in a Multi-GPU Architecture Using a Coalesced Data and Kernel Scheme. 546
Raphael Pereira Cordeiro, Rafael Sachetto Oliveira, Rodrigo Weber dos Santos, and Marcelo Lobosco

Service Dependability and Security in Distributed and Parallel Systems

Dynamic Verifiable Search Over Encrypted Data in Untrusted Clouds 557
Xiaohong Nie, Qin Liu, Xuhui Liu, Tao Peng, and Yapin Lin

Reducing TCB of Linux Kernel Using User-Space Device Driver 572
Weizhong Qiang, Kang Zhang, and Hai Jin

OBC Based Optimization of Re-encryption for Cryptographic Cloud Storage 586
Huidong Qiao, Jiangchun Ren, Zhiying Wang, Haihe Ba, Huaizhe Zhou, and Tie Hong

Performance Modeling and Evaluation

Modeling Performance of Hadoop Applications: A Journey from Queueing Networks to Stochastic Well Formed Nets 599
Danilo Ardagna, Simona Bernardi, Eugenio Gianniti, Soroush Karimian Aliabadi, Diego Perez-Palacin, and José Ignacio Requeno

D-SPACE4Cloud: A Design Tool for Big Data Applications 614
Michele Ciavotta, Eugenio Gianniti, and Danilo Ardagna

Porting MATLAB Applications to High-Performance C++ Codes: CPU/GPU-Accelerated Spherical Deconvolution of Diffusion MRI Data 630
Javier Garcia Blas, Manuel F. Dolz, J. Daniel Garcia, Jesus Carretero, Alessandro Daducci, Yasser Aleman, and Erick Jorge Canales-Rodriguez

On Stochastic Performance and Cost-Aware Optimal Capacity Planning of Unreliable Infrastructure-as-a-Service Cloud 644
Weiling Li, Lei Wu, Yunni Xia, Yuandou Wang, Kunyin Guo, Xin Luo, Mingwei Lin, and Wanbo Zheng

A Distributed Formal Model for the Analysis and Verification of Arbitration Protocols on MPSoCs Architecture 658
Imen Ben Hafaiedh, Maroua Ben Slimane, and Riadh Robbana

Synthetic Traffic Model of the Graph500 Communications 675
Pablo Fuentes, Enrique Vallejo, José Luis Bosque, Ramón Beivide, Andreea Anghel, Germán Rodríguez, Mitch Gusat, and Cyriel Minkenberg

Author Index 685