2013 20th International Conference on High Performance Computing

(HiPC 2013)

Bangalore, India
18-21 December 2013

Program

Conference Papers

Scheduling

[Analyzing the Performance Impact of Authorization Constraints and Optimizing the Authorization Methods for Workflows](#)

>   Nadeem Chaudhary (University of Warwick, United Kingdom); Ligang He (University of Warwick, United Kingdom)

[iFlatLFS: Performance Optimization for Accessing Massive Small Files](#)

>   Songling Fu (National University of Defense Technology, P.R. China); Ligang He (University of Warwick, United Kingdom)

[Adding Data Parallelism to Streaming Pipelines for Throughput Optimization](#)

>   Peng Li (Washington University, USA); Kunal Agrawal (Washington University in St. Louis, USA); Jeremy Daniel Buhler (Washington University, USA); Roger D Chamberlain (Washington University in St. Louis, USA)

[Algorithms for the Relaxed Multiple-Organization Multiple-Machine Scheduling Problem](#)

>   Anirudh Chakravorty (Indraprastha Institute of Information Technology, India); Neelima Gupta (University of Delhi, India); Neha Lawaria (University of Delhi, India); Pankaj Kumar (University of Delhi, India); Yogish Sabharwal (IBM Research - India, India)

[Loop Level Speculation in a Task Based Programming Model](#)

>   Rahulkumar Gayatri (Barcelona Supercomputing Center, Spain); Rosa M. Badia (Barcelona Supercomputing Center, Spain); Eduard Ayguade (Universitat Politècnica de Catalunya & Barcelona Supercomputing Center (BSC), Spain)

[EcoHadoop: A Cost-Efficient Data and Task Co-Scheduler for MapReduce](#)

Moussa Ehsan (Stony Brook University, USA); Yao Chen (Stony Brook University, USA); Hui Kang (Stony Brook University, USA); Radu Sion (Stony Brook University, USA); Jennifer L Wong (Stony Brook University, USA)

Virtualization & cloud computing

[Share-o-meter: An empirical analysis of KSM based memory sharing in virtualized systems](#)

Shashank Rachamalla (Indian Institute of Technology Bombay, India); Debadatta Mishra (Indian Institute of Technology Bombay, India); Purushottam Kulkarni (Indian Institute of Technology, Bombay, India)

[Minimization of Cloud Task Execution Length with Workload Prediction Errors](#)

Sheng Di (INRIA, France); Cho-Li Wang (The University of Hong Kong, Hong Kong)

[Speculative Dynamic Vectorization to Assist Static Vectorization in a HW/SW Co-designed Environment](#)

Rakesh Kumar (Universitat Politecnica de Catalunya, Barcelona, Spain); Alejandro Martínez (Intel Barcelona Research Center, Intel Labs, Spain); Antonio Gonzalez (Intel and UPC, Spain)

[A Self-Tuning System based on Application Profiling and Performance Analysis for Optimizing Hadoop MapReduce Cluster Configuration](#)

Dili Wu (Vanderbilt University, USA); Aniruddha Gokhale (Vanderbilt University, USA)

[Web-scale Entity Annotation Using MapReduce](#)

Shashank Gupta (IIT Bombay, India); Varun Chandramouli (NetApp India Pvt. Ltd., India); Soumen Chakrabarti (IIT Bombay, India)

[X10-Based Distributed and Parallel Betweenness Centrality and Its Application to Social Analytics](#)

Charuwat Houngkaew (Tokyo Institute of Technology, Japan); Toyotaro Suzumura (IBM Research - Tokyo & Tokyo Institute of Technology, Japan)

[Scheduling Associative Reductions with Homogeneous Costs when Overlapping Communications and Computations](#)

Louis-Claude Canon (Université de Franche-Comté & FEMTO-ST, France)

GPU computing

A Branch-and-Bound algorithm using multiple GPU-based LP solvers

> Paul Albuquerque (University of Applied Sciences of Western Switzerland, Switzerland); Xavier Meyer (University of Geneva, Switzerland); Bastien Chopard (University of Geneva, Switzerland)

Accelerating Strassen-Winograd's Matrix Multiplication Algorithm on GPUs

> Pai-Wei Lai (The Ohio State University, USA); Md Humayun Arafat (The Ohio State University, USA); Venmugil Elango (The Ohio State University, USA); Ponnuswamy Sadayappan (Ohio State University, USA)

Accelerating Inclusion-based Pointer Analysis on Heterogeneous CPU-GPU Systems

> Yu Su (University of New South Wales, Australia); Ding Ye (University of New South Wales, Australia); Jingling Xue (University of New South Wales, Australia)

Solving Tridiagonal Systems on a GPU

> Brian Murphy (Lehman College and The Graduate Center of the City University of New York, USA)

A Memory Efficient Algorithm for Adaptive Multidimensional Integration with Multiple GPUs

> Kamesh Arumugam (Old Dominion University, USA); Desh Ranjan (Old Dominion University, USA); Alexander Godunov (Old Dominion University, USA); Balsa Terzic (Jefferson Lab & Old Dominion University, USA); Mohammad Zubair (Old Dominion University, USA)

GAGM: Genome Assembly on GPU using Mate pairs

> Ashutosh Jain (Indian Institute of Technology, Delhi, India); Anshuj Garg (Indian Institute of Technology Delhi, India); Kolin Paul (Indian Institute of Technology, India)

Software

MaSiF: Machine Learning Guided Auto-tuning of Parallel Skeletons

> Alexander Collins (University of Edinburgh, United Kingdom); Christian Fensch (Heriot-Watt University, United Kingdom); Hugh Leather (University of Edinburgh, United Kingdom); Murray Cole (University of Edinburgh, United Kingdom)

[HARP: Adaptive Abort Recurrence Prediction for Hardware Transactional Memory](#)

Adrià Armejach (Barcelona Supercomputing Center & Universitat Politecnica de Catalunya, Spain); Anurag Negi (Chalmers University Of Technology, Sweden); Osman Unsal (Barcelona Supercomputing Center, Spain); Adrian Cristal (Barcelona Supercomputing Center, Spain); Per Stenstrom (Chalmers University of Technology, Sweden); Timothy Harris (Oracle Labs, Cambridge, Germany)

[MIL: A language to build program analysis tools through static binary instrumentation](#)

Andres S. Charif-Rubial (University of Versailles Saint-Quentin en Yvelines, France); Denis Barthou (University of Bordeaux - Labri / INRIA, France); Cédric Valensi (University of Versailles St Quentin, France); Sameer S Shende (University of Oregon & ParaTools, Inc., USA); Allen D. Malony (University of Oregon, USA); William Jalby (University of Versailles Saint-Quentin en Yvelines, France)

[Cache-Based Cross-Iteration Coherence for Speculative Parallelization](#)

Andre Baixo (University of Washington, Brazil); Joao Porto (Google Inc., USA); Guido Araujo (University of Campinas (UNICAMP), Brazil)

[Exploring Energy and Performance Behaviors of Data-Intensive Scientific Workflows on Systems with Deep Memory Hierarchies](#)

Marc Gamell (Rutgers University, USA); Ivan Rodero (Rutgers University & NSF CAC, USA); Manish Parashar (Rutgers, The State University of New Jersey, USA); Stephen W. Poole (Oak Ridge National Laboratory, USA)

[Transaction Scheduling using Conflict Avoidance and Contention Intensity](#)

Marcio Pereira (University of Campinas - UNICAMP, Brazil); Alexandro Baldassin (Universidade Estadual Paulista, Brazil); Luiz Buzato (Universidade Estadual de Campinas, Brazil); Guido Araújo (University of Campinas - UNICAMP, Brazil)

Algorithms

[A New Parallel Algorithm for Connected Components in Dynamic Graphs](#)

Robert McColl (Georgia Institute of Technology & Georgia Tech Research Institute, USA); Oded Green (Georgia Institute of Technology, USA); David A. Bader (Georgia Institute of Technology, USA)

[The Super Warp Architecture with Random Address Shift](#)

Koji Nakano (Hiroshima University, Japan); Susumu Matsumae (Saga University, Japan)

[Parallel Branch-and-Bound for Two-Stage Stochastic Integer Optimization](#)

Akhil Langer (University of Illinois at Urbana-Champaign, USA); Ramprasad Venkataraman (University of Illinois, USA); Udatta Palekar (University of Illinois at Urbana-Champaign, USA); Laxmikant V. Kale (University of Illinois at Urbana-Champaign, USA)

[A Dynamic Schema to increase performance in Many-core Architectures through Percolation operations](#)

Elkin Garcia (University of Delaware, USA); Daniel A Orozco (University of Delaware, USA); Rishi Khan (ET International, USA); Ioannis E. Venetis (University of Patras, Greece); Kelly Livingston (University of Delaware, USA); Guang Gao (University of Delaware, USA)

[Efficient sparse matrix multiple-vector multiplication using a bitmapped format](#)

Ramaseshan Kannan (University of Manchester, United Kingdom)

[Approximation Algorithms for Energy Minimization in Cloud Service Allocation under Reliability Constraints](#)

Olivier Beaumont (Inria, France); Philippe Duchon (LaBRI, university of Bordeaux 1, France); Paul RENAUD-GOUD (U. of Bordeaux, France)

GPU and hybrid computing

[Can GPUs Sort Strings Efficiently??](#)

Aditya Deshpande (International Institute of Information Technology, India); P J Narayanan (International Institute of Information Technology, India)

[Parallel Distributed Breadth First Search on GPU](#)

Koji Ueno (Tokyo Institute of Technology, Japan); Toyotaro Suzumura (Tokyo Institute of Technology & IBM Research - Tokyo, Japan)

[Evaluation and Enhancement of Weather Application Performance on Blue Gene/Q](#)

Gurbinder Gill (IBM Research-India, New Delhi, India); Vaibhav Saxena (IBM Research - India, New Delhi, India); Rashmi Mittal (India & IBM Research, India); Thomas George (IBM Research India, India); Yogish Sabharwal (IBM Research - India, India); Lalit Dagar (Universiti Brunei Darussalam, Brunei Darussalam)

[Efficient Homology Computations on Multicore and Manycore Systems](#)

Anurag Murty (Indian Institute of Science, India); Vijay Natarajan (Indian Institute of Science Bangalore India, India); Sathish Vadhiyar (Indian Institute of Science, India)

[A hybrid shared memory heterogeneous execution platform for PCIe-based GPGPUs](#)

Sambit Shukla (University of California, Riverside, USA); Laxmi Bhuyan (University of California, USA)

[GPU-enabled Efficient Executions of Radiation Calculations in Climate Modeling](#)

Sai Kiran Korwar (Indian Institute of Science, India); Sathish Vadhiyar (Indian Institute of Science, India); Ravi S Nanjundiah (Indian Institute of Science, India)

Applications

[Revisiting the space-filling curves for storage, reordering and partitioning mesh based data in scientific computing](#)

Pavanakumar Mohanamuraly (CSIR National Aerospace Laboratories, India); Kaushik Kumar Nagarajan (CSIR National Aerospace Laboratories, India)

[Multi tier energy buffering management for IDCs with heterogeneous energy storage devices](#)

Zahra Abbasi (Arizona State University, USA); Madhurima Pore (Arizona State University, USA); Ayan Banerjee (Arizona State University, USA); Sandeep Gupta (Arizona State University, USA)

[Benchmarking MIC architectures with Monte Carlo simulations of spin glass systems](#)

Alessandro Gabbana (University of Ferrara, Italy); Marcello Pivanti (University of Ferrara, Italy); Sebastiano Fabio Schifano (University of Ferrara & INFN, Italy); Raffaele Tripiccione (University of Ferrara, Italy)

[Performance and Energy Consumption Analysis of a Seismic Application for Three different Architectures intended for Oil and Gas Industry](#)

Lucas Melo (Federal University of Pernambuco, Brazil); Gilliano Menezes (Federal University of Pernambuco, Brazil); Abel Silva-Filho (Federal University of Pernambuco, Brazil); Manoel Lima (UFPE, Brazil)

[Performance Evaluation of Medical Imaging Algorithms on Intel MIC Platform](#)

Jyotsna Khemka (Siemens Corporate Reseach and Technology & Siemens AG, India); Mrugesh Gajjar (Siemens Corporate Research, India); Sharan Vaswani (Siemens Corporate Research and Technologies, India); Nagavijayalakshmi Vydyanathan (Siemens Corporate Technology, India); Rama Malladi (Intel, India); Vinutha V (Intel, India)

[A Hybrid Parallelization Approach for High Resolution Operational Flood Forecasting](#)

Swati Singhal (IBM India Research Lab, India); Lucas Villa Real (IBM Research Brazil, Brazil); Thomas George (IBM Research India, India); Sandhya Aneja (University of Delhi South Campus, India); Yogish Sabharwal (IBM Research - India, India)

Algorithms and networking

[Effects of Phase Imbalance on Data Center Energy Management](#)

Sushil Gupta (HCL Infosystems Ltd, India); Ayan Banerjee (Arizona State University, USA); Zahra Abbasi (Arizona State University, USA); Sandeep Gupta (Arizona State University, USA)

[Conflict-free data access for multi-bank memory architectures using padding](#)

Joar Sohl (Linköping University, Sweden); Jian Wang (Linköping University, Sweden); Andréas Karlsson (Linköping University, Sweden); Dake Liu (Linköping University, Sweden)

[Work Efficient Parallel Algorithms for Large Graph Exploration](#)

Dip Sankar Banerjee (International Institute Of Information Technology, Hyderabad, India); Shashank Sharma (IIIT Hyderabad, India); Kishore Kothapalli (International Institute of Information Technology, India)

[SCORPIO: A Scalable Two-Phase Parallel I/O Library With Application To A Large Scale Subsurface Simulator](#)

Sarat Sreepathi (Oak Ridge National Laboratory & Computer Science and Mathematics Division, USA); Vamsi Sripathi (Intel Corporation, USA); Richard Mills (Oak Ridge National Laboratory, USA); Glenn Hammond (Pacific Northwest National Laboratory, USA); G. Kumar Mahinthakumar (North Carolina State University, USA)

[Compiler Generation and Autotuning of Communication-Avoiding Operators for Geometric Multigrid](#)

Protonu Basu (University of utah, USA); Samuel W. Williams (Lawrence Berkeley National Laboratory, USA); Brian Van Straalen (Lawrence Berkeley National Laboratory, USA); Anand Venkat (University of Utah, USA); Leonid Oliker (Lawrence Berkeley National Laboratory, USA); Mary Hall (University of Utah, USA)

SymSig: A Low Latency interconnection topology for HPC clusters

Dhananjay Brahme (Tata Consultancy Services, India); Onkar Bhardwaj (Rensselaer Polytechnic Institute, USA); Vipin Chaudhary (University at Buffalo, SUNY, USA)

HiPC-13 Workshops Proceedings ¨B#5

International Workshop on Cloud Computing Applications (IWCA)

Composite Prediction Model and Task Distribution on a Cloud of Multi-core Prcessors

Khondker S. Hasan (University of Oklahoma, USA); Sridhar Radhakrishnan (University of Oklahoma, USA); John K. Antonio(University of Oklahoma, USA)

Efficient Support of Big Data Storage Systems on the Cloud

Akshay MS (PES Institute of Technology, India); Suhas Mohan (PES Institute of Technology, India) Vincent Kuri (PES Institute of Technology, India) Dinkar Sitaram, H. (PES Institute of Technology, India) L. Phalachandra (PES Institute of Technology, India)

Proffer the Legacy System to Cloud Environment: An Approach

Ajeet A. Chikkamannur (Sri Venkateshwara College of Engineering, India); Shivanand M. Handigund (Sri Venkateshwara College of Engineering, India)

Workshop on the Performance Engineering and Applications (WPEA)

A Data-Parallel Implementation of Quadrature Methods for Complex Numbers

Sandeep Koranne (Mentor Graphics Corporation, USA)

A Novel Technique to Improve Parallel Program Performance Co-executing with Dynamic Workloads

Murali Krishna Emani (School of Informatics, University of Edinburgh, UK); Michale O'Boyle(School of Informatics, University of Edinburgh, UK)

Performance trends of multicore system for throughput

Madhurima Pore (Arizona State University, USA); Ayan Banerjee (Arizona State University, USA); Sandeep K. S. Gupta (Arizona State University, USA); Hari K Tadepalli (Intel, USA)

Block Dimension Selection for GPU Kernels using Artificial Neural Networks

5f_U; \cg\ fGf]´GUh\nŮGU]´=bgh]li hŸcZ<][\Yf´@YŮfb]b[ž=bX]Ł/´G¨
6UŮgj VfŮa Ŭb]ŬbfGf]´GUh\nŮGU]´=bgh]li hŸcZ<][\Yf´@YŮfb]b[ž=bX]Ł