

# **2014 21st International Conference on High Performance Computing**

**(HiPC 2014)**

**Goa, India  
17-20 December 2014**



**IEEE Catalog Number: CFP14176-POD  
ISBN: 978-1-4799-5977-8**

Time            Session

## Thursday, December 18

10:00-12:00 [Network and Computer Architectures](#)

13:30-15:30 [MPI, I/O and Interconnects](#)

16:00-18:00 [Graph Partitioning and Numerical Algorithms](#)

## Friday, December 19

10:00-12:00 [Many Core and Storage](#)

13:30-15:30 [Scheduling and Cloud Computing](#)

16:00-18:00 [Applications: GPU Computing](#)

## Saturday, December 20

10:00-12:00 [Software for HPC and Distributed Computing](#)

13:30-16:00 [Algorithms and Computing on Accelerators](#)

---

## Thursday, December 18

10:00 - 12:00

### Network and Computer Architectures

[Optimizing the performance of parallel applications on a 5D torus via task mapping](#)\*\*\*\*%

Abhinav Bhatele (Lawrence Livermore National Laboratory, USA); Nikhil Jain (University of Illinois at Urbana-Champaign, USA); Katherine Isaacs (University of California, Davis, USA); Ronak Buch (University of Illinois at Urbana-Champaign, USA); Todd Gamblin and Steven Langer (Lawrence Livermore National Laboratory, USA); Laxmikant V. Kale (University of Illinois at Urbana-Champaign, USA)

[Balancing Context Switch Penalty and Response Time with Elastic Time Slicing](#)\*\*\*\*%

Nagakishore Jammula, Moinuddin Qureshi, Ada Gavrilovska and Jongman Kim (Georgia Institute of Technology, USA)

[Scaling Graph Community Detection on the Tiler Many-core Architecture](#)\*\*\*\*&

Daniel Gerardo Chavarria and Mahantesh Halappanavar (Pacific Northwest National Laboratory, USA); Ananth Kalyanaraman (Washington State University, USA)

[Matrix-Matrix Multiplication on a Large Register File Architecture with Indirection](#)\*\*\*\*' &

Dheeraj Sreedhar (IBM India Research Labs, India); Jeff H. Derby (IBM, USA); Robert Montoye and Charles Johnson (IBM Research Division, USA)

[Trikon: A Hypervisor Aware Manycore Processor](#)\*\*\*\*' (&

Rohan Bhalla and Prathmesh Kallurkar (Indian Institute of Technology, Delhi, India); Nitin Gupta (Amazon India pvt ltd, India); Smruti Sarangi (IIT Delhi, India)

[Optical Overlay NUCA: A High Speed Substrate for Shared L2 Caches](#) &

Eldhose Peter (Indian Institute of Technology Delhi, India); Anuj Arora (Indian Institute of Technology, Delhi, India); Akriti Bagaria and Smruti Sarangi (IIT Delhi, India)

**13:30 - 15:30**

**MPI, I/O and Interconnects**

[On the Suitability of MPI as a PGAS Runtime](#) &

Jeff Daily (Pacific Northwest National Laboratory & Washington State University, USA); Abhinav Vishnu and Bruce Palmer (Pacific Northwest National Laboratory, USA); Hubertus J. J. Van Dam (Pacific Northwest National Lab, USA); Darren J. Kerbyson (Pacific Northwest National Laboratory, USA)

[Designing Efficient Small Message Transfer Mechanism for Inter-node MPI Communication on InfiniBand GPU Clusters](#) +&

Rong Shi, Sreeram Potluri, Khaled Hamidouche, Jonathan Perkins and Mingzhe Li (The Ohio State University, USA); Davide Rossetti (NVIDIA Corporation, USA); Dhabaleswar Panda (The Ohio State University, USA)

[Combining HoL-blocking Avoidance and Differentiated Services in High-Speed Interconnects](#) &

Pedro Yebenes, Jesús Escudero-Sahuquillo, Crispin Gomez, Pedro Garcia and Francisco Jose Alfaro (University of Castilla-La Mancha, Spain); Francisco J. Quiles (Universidad Castilla La Mancha, Spain); Jose Duato (Universidad Politecnica de Valencia, Spain)

[A High Performance Broadcast Design with Hardware Multicast and GPUDirect RDMA for Streaming Applications on Infiniband Clusters](#) - &

Akshay Venkatesh (Ohio State University, USA); Hari Subramoni (Senior Research Associate, USA); Khaled Hamidouche and Dhabaleswar Panda (The Ohio State University, USA)

[High Performance MPI Library over SR-IOV Enabled InfiniBand Clusters](#) %&

Jie Zhang, Xiaoyi Lu, Jithin Jose, Mingzhe Li, Rong Shi and Dhabaleswar Panda (The Ohio State University, USA)

[DRIVE: Using Implicit Caching Hints to achieve Disk I/O Reduction in Virtualized Environments](#) %&

Sujesha Sudevalayam (IIT Bombay, India); Purushottam Kulkarni (Indian Institute of Technology, Bombay, India); Rahul Balani (IBM Research India, India); Akshat Verma (IBM Research - India, India)

**16:00 - 18:00**

**Graph Partitioning and Numerical Algorithms**

[An improved recursive graph bipartitioning algorithm for well balanced domain decomposition](#) %&&

Astrid Casadei (Bordeaux University & Inria, France); Pierre Ramet (University of Bordeaux & INRIA - LaBRI, France); Jean Roman (INRIA, France)

[Coupling-Aware Graph Partitioning Algorithms: Preliminary Study](#) % &

Aurélien Esnard (University of Bordeaux, LaBRI, INRIA Bordeaux Sud-Ouest, HiePACS, France); Maria Predari (University of Bordeaux, INRIA Bordeaux Sud-Ouest, HiePACS, France)

[Reducing elimination tree height for parallel LU factorization of sparse unsymmetric matrices](#) % &

Enver Kayaaslan (INRIA, France); Bora Uçar (CNRS & ENS Lyon, France)

[Analysis and Tuning of Libtensor Framework on Multicore Architectures](#) % &

Khaled Z Ibrahim (Lawrence Berkeley National Laboratory, USA); Samuel W. Williams (Lawrence Berkeley National Laboratory, USA); Evgeny Epifanovskiy (UC Berkeley, USA); Anna Krylov (University of Southern California, USA)

[A Multilevel Compressed Sparse Row Format for Efficient Sparse Computations on Multicore Processors](#) % &

Humayun Kabir, Joshua Booth and Padma Raghavan (The Pennsylvania State University, USA)

[Optimization of Scan Algorithms on Multi- and Many-core Processors](#) %&

Qiao Sun and Chao Yang (Institute of Software, Chinese Academy of Sciences, P.R. China)

---

# Friday, December 19

10:00 - 12:00

## Many Core and Storage

[Heterogeneous many cores for medical control: Performance, Scalability, and Accuracy](#) &

Madhurima Pore, Ayan Banerjee and Sandeep Gupta (Arizona State University, USA)

[Interface for Heterogeneous Kernels: A Framework to Enable Hybrid OS Designs targeting High Performance Computing on Manycore Architectures](#) &

Taku Shimosawa (Hitachi, Ltd., Japan); Balazs Gerofi (The University of Tokyo, Japan); Masamichi Takagi (RIKEN AICS, Japan); Gou Nakamura (Hitachi Solutions, Ltd., Japan); Tomoki Shirasawa (Hitachi Solutions East Japan, Ltd., Japan); Yuji Saeki (Hitachi, Ltd., Japan); Masaaki Shimizu (Central Research Laboratory, Hitachi, Ltd., Japan); Atsushi Hori (Riken & Advanced Institute for Computer Science, Japan); Yutaka Ishikawa (University of Tokyo, Japan)

[Premonition of Storage Response Class Using Skyline Ranked Ensemble Method](#) &&

Kumar Dheenadayalan (Qualcomm India Pvt Ltd, India); Muralidhara V n (IIIT Bangalore, India); Pushpa Datla (Qualcomm, India); G Srinivasaraghavan (IIITB, India); Maulik Shah (Qualcomm, India)

[Queueing-based Storage Performance Modeling and Placement in OpenStack Environments](#) &&

Yang Song and Rakesh Jain (IBM Research, USA); Ramani Routray (IBM Almaden Research Center, USA)

[A Fast Implementation of MLR-MCL Algorithm on Multi-core Processors](#) &&&

Qingpeng Niu, Pai-Wei Lai, SM Faisal and Srinivasan Parthasarathy (The Ohio State University, USA); Ponnuswamy Sadayappan (Ohio State University, USA)

[Optimizing Shared Data Accesses in Distributed-Memory X10 Systems](#) & &

Jeeva Paudel (University of Alberta, Canada); Olivier Tardieu (IBM Research, USA); Jose Nelson Amaral (University of Alberta, Canada)

13:30 - 15:30

## Scheduling and Cloud Computing

[A Proactive Approach for Coping with Uncertain Resource Availabilities on Desktop Grids](#) & (&

Louis-Claude Canon (Université de Franche-Comté & FEMTO-ST, France); Adel Essafi (University of Grenoble, France); Denis R. Trystram (Grenoble Institute of Technology, France)

[Algorithms for Power-Aware Resource Activation](#) & & %

Archita Agarwal (IBM Research - India, India); Sonika Arora (University of Delhi, India); Venkatesan T Chakaravarthy (IBM Research (India), India); Yogish Sabharwal (IBM Research - India, India)

[A Flexible Scheduling Framework for Heterogeneous CPU-GPU Clusters](#) & \* %

Kittisak Sajjapongse and Tejaswi Agarwal (University of Missouri, USA); Michela Becchi (University of Missouri - Columbia, USA)

[Cache-Conscious Scheduling of Streaming Pipelines on Parallel Machines with Private Caches](#) & + &

Kunal Agrawal (Washington University in St. Louis, USA); Jeremy Fineman (Georgetown University, USA); Jordyn Maglalang (Washington University in St. Louis, USA)

[Efficient and Robust Allocation Algorithms in Clouds under Memory Constraints](#) & (

Olivier Beaumont (Inria, France); Lionel Eyraud-Dubois (INRIA Bordeaux Sud-Ouest & University of Bordeaux,, France); Juan-Angel Lorenzo (Inria, France); Paul Renaud-Goud (Chalmers University of Technology, Sweden)

[Saving Energy by Exploiting Residual Imbalances on Iterative Applications](#) & - (

Edson Luiz Padoin (Federal University of Rio Grande do Sul and Laboratoire d'Informatique de Grenoble, Brazil); Márcio Castro (Federal University of Rio Grande do Sul (UFRGS), Brazil); Laercio L. Pilla (UFSC, Brazil); Philippe O. A. Navaux (Universidade Federal do Rio Grande do Sul & Informatic Institute, Brazil); Jean-François Méhaut (Grenoble University, France)

**16:00 - 18:00**

### **Applications: GPU Computing**

[GPU Parallelization of the Stochastic On-time Arrival Problem](#) ( ) \$ (

Maleen Abeydeera (University of Moratuwa, Sri Lanka); Samitha Samaranayake (UC Berkeley, USA)

[GpuTejas: A Parallel Simulator for GPU Architectures](#) ( ) %&

Geetika Malhotra and Seep Goel (Indian Institute of Technology, Delhi, India); Smruti Sarangi (IIT Delhi, India)

[Mixed-Precision Models for Calculation of High-Order Virial Coefficients on GPUs](#) ( ) &&

Chao Feng, Andrew Schultz, Vipin Chaudhary and David Kofke (University at Buffalo, SUNY, USA)

[Parallel AMG Solver for Three Dimensional Unstructured Grids Using GPU](#) ( ) &

RaviTej Kamakolanu (Indian Institute of Technology Hyderabad, India); Naveen Sivadasan (IITH, India); Vatsalya Sharma (Indian Institute of Technology, India); Raja Banerjee (Indian Institute of Technology Hyderabad, India)

[Particle Advection Performance Over Varied Architectures and Workloads](#) ( ) ( &

Hank Childs (LBL, USA); Scott Biersdorff and David Poliakov (University of Oregon, USA); David Camp (Lawrence Berkeley National Laboratory, USA); Allen D. Malony (University of Oregon, USA)

[Relax-Miracle: GPU Parallelization of Semi-Analytic Fourier-Domain solvers for Earthquake Modeling](#) ( ) &

Nachiket Kapre (Nanyang Technological University & Imperial College London, Singapore)

---

## **Saturday, December 20**

**10:00 - 12:00**

### **Software for HPC and Distributed Computing**

[Xevolver: An XML-based Code Translation Framework for Supporting HPC Application Migration](#) ( ) \* &

Hiroyuki Takizawa (Tohoku University, Japan); Shoichi Hirasawa (TOHOKU University, Japan); Yasuharu Hayashi (NEC Corporation, Japan); Ryusuke Egawa (Tohoku University Japan, Japan); Hiroaki Kobayashi (Tohoku University, Japan)

[Online failure prediction for HPC resources using decentralized clustering](#) ( ) +'

Alejandro Pelaez (Rutgers University, USA); Manish Parashar (Rutgers, The State University of New Jersey, USA); Andres Quiroz (Xerox Research, USA); Jim Browne (UT Austin, USA); Edward Chuah (University of Texas at Austin, USA)

[COA: A Code Quality Analyzer tool at binary level](#) ( ) , &

Andres S. Charif-Rubial and Emmanuel Oseret (University of Versailles Saint-Quentin en Yvelines, France); Ghislain Lartigue (University of Normandie, France); Jose Noudohouenou and William Jalby (University of Versailles Saint-Quentin en Yvelines, France)

[Towards Realizing the Potential of Malleable Jobs](#) ( ) - &

Abhishek Gupta (Intel Corp, USA); Bilge Acun (University of Illinois at Urbana Champaign, USA); Osman Sarood (University of Illinois Urbana Champaign, USA); Laxmikant V. Kale (University of Illinois at Urbana-Champaign, USA)

[Improving Multi-dimensional Query Processing with Data Migration in Distributed Cache Infrastructure](#) ( ) (\$&

Youngmoon Eom (Ulsan National Institute of Science and Technology, Korea); Jinwoong Kim, Deukyeon Hwang and Jaewon Kwak (UNIST, Korea); Minh Shin (Myongji University, Korea); Beomseok Nam (Ulsan National Institute of Science and Technology, Korea)

[An Early Experience of Regional Ocean Modelling on Intel Many Integrated Core Architecture](#) ( ) (%&

Srikanth Yalavarthi and Akshara Kaginalkar (Centre for Development of Advanced Computing, India)

**13:30 - 16:00**

### **Algorithms and Computing on Accelerators**

- [RADIR: Lock-free and Wait-free Bandwidth Allocation Models for Solid State Drives](#) (%, Pooja Aggarwal (Indian Institute of Technology, Delhi, India); Giridhar Yasa (NetApp, India); Smruti Sarangi (IIT Delhi, India))
- [Design and Evaluation of Parallel Hashing over Large-scale Data](#) (&, Long Cheng (Technische Universität Dresden, Germany); Spyros Kotoulas (IBM Research, Ireland); Tomas Ward (National University of Ireland Maynooth, Ireland); Georgios Theodoropoulos (Durham University, United Kingdom))
- [Smart Multi-Task Scheduling for OpenCL Programs on CPU/GPU Heterogeneous Platforms](#) (' , Yuan Wen (University of Edinburgh, United Kingdom); Zheng Wang (Lancaster University, United Kingdom); Michael O'Boyle (Edinburgh, United Kingdom))
- [Software Based Ultrasound B-mode/Beamforming Optimization on GPU and its Performance Prediction](#) ((, Thi Yen Phuong and Jeong-Gun Lee (Hallym University, Korea))
- [Fine-grained GPU parallelization of Pairwise Local Sequence Alignment](#) ((, Chirag Jain (Indian Institute of Technology, Delhi, India); Subodh Kumar (IIT Delhi, India))
- [Distance Threshold Similarity Searches on Spatiotemporal Trajectories using GPGPU](#) ((\*, Michael Gowanlock and Henri Casanova (University of Hawaii at Manoa, USA))
- [Simple Parallel Biconnectivity Algorithms for Multicore Platforms](#) ((+, George Slota and Kamesh Madduri (The Pennsylvania State University, USA))