24th IEEE International Conference on High Performance Computing, Data, and Analytics

Jaipur, India • December 18-21, 2017

PROGRAM OVERVIEW

Starting on Monday, December 18th, the first day of HiPC 2017 features two workshops, two tutorials, and a Birds-of-a-Feather session on gem5 and concludes with an Industry Gala. The technical program on Days 2, 3 and 4 will showcase three keynote speakers and seven single track sessions of forty-one peer reviewed papers from all over the world covering important and timely topics in all areas of high performance computing, data, and analytics. Also on Days 2 and 3, the conference will present plenary panels covering topics of special interest to the HiPC community.

On Days 2 and 3, there will be a full program of industry exhibits and related events including Industry, Research and Users Symposium (IRUS) sessions to provide a forum for presenting state-of-the-art in HPC platforms and technologies, discussing best practices, and exchanging experiences. In addition, industry representatives will host special technical sessions on these two days. On Days 2 and 3, posters of work selected for the Student Research Symposium will be on display.

KEYNOTE SPEAKERS

Tuesday, December 19th – Day 2 Keynote

End of Moore's Law: Or, a computer architect's mid-life crisis?

Partha Ranganathan
Google

Wednesday, December 20th – Day 3 Keynote

Machine Learning @ Amazon

Rajeev Rastogi
Amazon India

Thursday, December 21st – Day 4 Keynote

Computing Just What You Need: Online Data Analysis and Reduction at Extreme Scales

Ian Foster
Argonne National Lab and University of Chicago

TECHNICAL PROGRAM

The HiPC 2017 technical program on Days 2, 3 and 4 will consist of 41 peer reviewed papers chosen from nearly 200 submissions from all over the world. The papers will be presented in seven single-track sessions and will cover important and timely topics in all areas of high performance computing and data science. The three keynote speakers will open the morning plenary sessions. On Days 2 and 3, two plenary panels will be held at the end of the day, one on HPC initiatives in India, the second on machine learning and mathematical modeling.

WORKSHOPS

Two workshops covering diverse topics complementary to the conference technical program will be held on Day 1 of the conference.

1. Foundations of Big Data Computing (BiGDF)
2. Computational Fluid Dynamics (CFD)

Peer reviewed papers selected for presentation in these workshops are included in the proceedings. Abstracts for keynote speakers and invited speakers are included in the Workshop Introduction file in the proceedings as well as the individual webpages for the workshops.

STUDENT RESEARCH SYMPOSIUM

The 10th HiPC Student Research Symposium is aimed at stimulating and fostering student research, and providing an international forum to highlight student research accomplishments in HPC. The symposium also gives students exposure to the best practices of senior HPC researchers in academia and industry. As a program feature, started last year, the 2017 symposium will feature only student posters - there will be no talks by students – and the posters will be on display on both Day 2 and Day 3 of the conference, next to industry exhibits. This will allow students attending to more fully participate in all events of the conference.

ACADEMIC BoF SESSIONS

HiPC introduced this format in 2014 with the purpose of fostering greater participation by the academic community at HiPC, including both domestic and international representation. With a focus on emerging areas of interest, the goal is to stimulate new research ideas, address specific problems, and build a research community. This year, one will be held on Day 1 and another on Day4.

INDUSTRY EVENTS

The conference welcomes (and strongly encourages) industry participation on all days at all levels including in the technical program and student symposium. The industry/research exhibition, to be held on December 19th and 20th, will include booths and demonstrations and will showcase products, services and current work from vendor companies and R&D laboratories. **Two Industry, Research and Users Symposium (IRUS) sessions** will be held on Day 2 and will bring together solution providers and users of HPC in a forum to discuss platforms and technologies and best practices. **Three lead industry sponsors will conduct technical sessions: SanDisk and Microsoft on Day 2, and Xilinx on Day 3.**

HiPC 2017 Venue and Meeting Rooms

Le Méridien
Jaipur, India
Number 1, RIICO, Kukas
Jaipur, Rajasthan 302028 India
Phone: (91)(1426) 515100

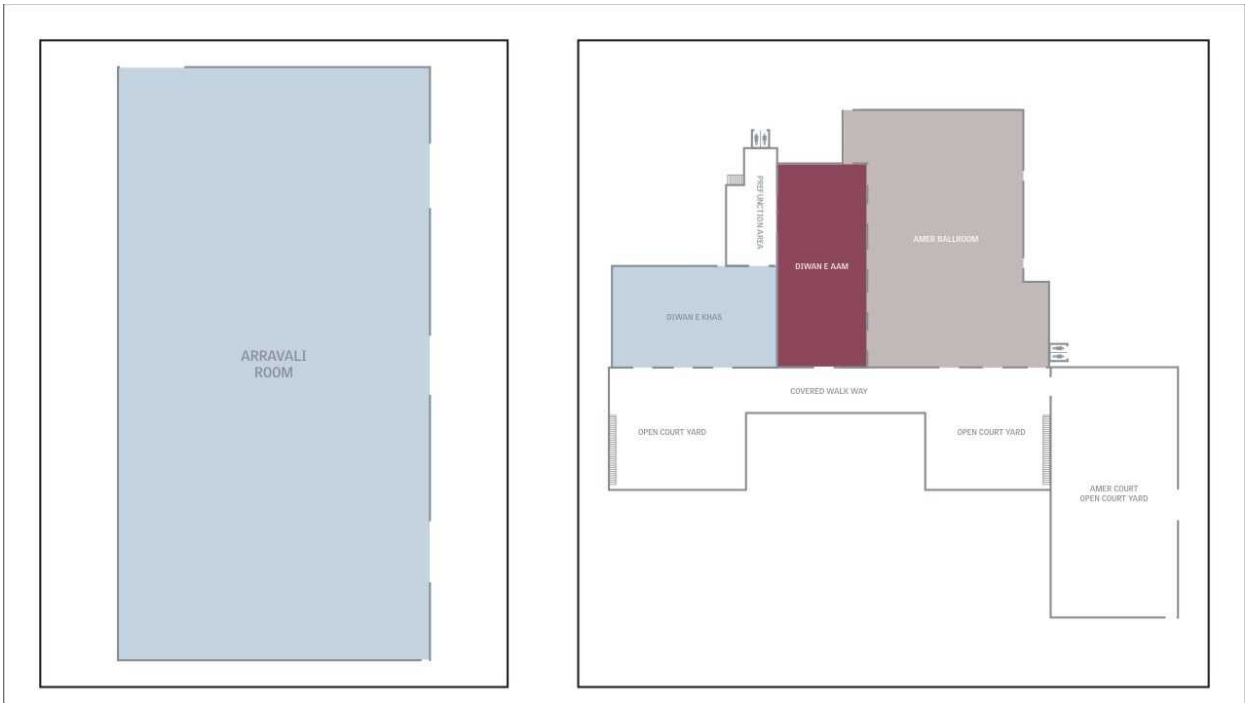

Aravali • Diwan E Khas • Diwan E Aam • Amer Ballroom

Royal Lounge • Surya Vilas

HiPC 2017 PROGRAM DAILY SCHEDULE

Check the Mobile App and postings at the venue for up-to-date information on locations and schedules for each event. The following shows time and location for all events.

Day	Start Time	End Time	Event	Le Méridien Room
Mon			*Breakfast @ 7:30 – Lunch @ 1:00*	Refreshment Lounge
Mon			*Breaks: Morning Break@ 10:30 – Afternoon @ 4:00*	Lobby Area
Mon	8:30 AM	6:30 PM	**Industry Exhibit Booth Set-up**	Diwan E Aam
Mon	8:30 AM	6:30 PM	**Workshop 1: BIGDF**	Diwan E Khas
Mon	8:30 AM	6:30 PM	**Workshop 2: CFD**	Aravali Hall
Mon	11:00 AM	1:00 PM	**Academic BOF1: GEM5**	Surya Vilas
Mon	2:00 PM	4:00 PM	**Tutorial 1: Power Optimization**	Surya Vilas
Mon	4:30 PM	6:30 PM	**Tutorial 2: FPGAs for HPC**	Surya Vilas
Mon	4:30 PM	6:30 PM	**Student Research Symposium Poster Set-up**	Royal Lounge
Mon	6:30 PM	8:00 PM	**Plenary Industry Gala Opening**	TBA
Tue			*Breakfast @ 7:00 – Lunch @ 12:00*	Refreshment Lounge
Tue			*Breaks: Morning @ 9:30 – Afternoon @ 3:00 – Late @ 5:15 PM*	Lobby Area
Tue	8:00 AM	8:30 AM	**Plenary HiPC 2017 Inauguration**	Amer Ballroom
Tue	8:30 AM	9:30 AM	**Keynote Talk 1: Parthasarathy Ranganathan (Google)**	Amer Ballroom
Tue	10:00 AM	12:00 PM	**Tech Session 1: Graph Algorithms**	Amer Ballroom
Tue	10:00 AM	12:00 PM	**IRUS 1: Applications of Machine Learning in e-commerce**	Aravali Hall
Tue	10:00 AM	12:00 PM	**Sponsored Tech Session: SanDisk**	Diwan E Khas
Tue	10:00 AM	5:30 PM	**Industry Demos & Exhibits**	Diwan E Aam
Tue	10:00 AM	5:30 PM	**SRS Posters on display**	Royal Lounge
Tue	1:00 PM	3:00 PM	**Tech Session 2: Architecture and Communication**	Amer Ballroom
Tue	1:00 PM	3:00 PM	**Sponsored Tech Session: Microsoft**	Diwan E Khas
Tue	3:15 PM	5:15 PM	**Tech Session 3: Algorithms**	Amer Ballroom
Tue	3:15 PM	5:15 PM	**IRUS 2: Performance Challenge in Deep Learning and Scientific Computing**	Aravali Hall
Tue	5:30 PM	7:30 PM	**Plenary Panel: HPC in India**	Amer Ballroom
Wed			*Breakfast @ 7:30 – Lunch @ 12:00*	Refreshment Lounge
Wed			*Breaks: Morning Break@ 9:30 – Afternoon @ 3:00*	Lobby Area
Wed	8:30 AM	9:30 AM	**Keynote Talk 2: Rajeev Rastogi (Amazon)**	Amer Ballroom
Wed	10:00 AM	12:00 PM	**Tech Session 4: Big Data, Machine Learning and Optimization**	Amer Ballroom
Wed	10:00 AM	12:00 PM	**Faculty Summit**	Surya Vilas
Wed	10:00 AM	5:30 PM	**Industry Demos & Exhibits**	Diwan E Aam
Wed	10:00 AM	5:30 PM	**SRS Posters on display**	Royal Lounge
Wed	1:00 PM	3:00 PM	**Tech Session 5: Graph Algorithms and GPU**	Amer Ballroom
Wed	1:00 PM	3:00 PM	**Sponsored Tech Session: Xilinx**	Diwan E Khas
Wed	3:30 PM	5:30 PM	**Plenary Panel: Machine Learning and Mathematical Modeling**	Diwan E Khas
Wed	6:30 PM	9:30 PM	**Banquet & Awards**	Amer Ballroom
Thu			*Breakfast @ 7:30 – Lunch @ 12:00*	Refreshment Lounge
Thu			*Morning Break@ 9:30*	Lobby Area
Thu	8:30 AM	9:30 AM	**Keynote Talk 3: Ian Foster (Argonne Nat Lab & U of Chicago)**	Amer Ballroom
Thu	10:00 AM	12:00 PM	**Tech Session 6: System Software**	Amer Ballroom
Thu	10:00 AM	12:00 PM	**Academic BOF: Parallel LSG Computation**	Aravali Hall
Thu	1:00 PM	3:00 PM	**Tech Session 7: GPU Frameworks and Applications**	Amer Ballroom

Day 1 - Monday, December 18

Breakfast @ 7:30, Breaks @ 10:30 and 4:00, Lunch @ 1:00

HiPC 2017 WORKSHOP 1:
Foundations of Big Data Computing (BIGDF)

Organizers:
Madhu Govindaraju, Saumyadipta Pyne, Dinkar Sitaram, Ananth Kalyanaraman

08:50-10:30
Session 1
Opening Remarks by Workshop Chairs

Software Troubleshooting using Machine Learning
Neha M Kalibhat, Shreya Varshini, Chidambaram Kollengode, Dinkar Sitaram, Subramaniam Kalambur

The Effect of Non Volatile Memory on a Distributed Storage System
Ramdoot Pydipaty, Amit Saha, Johnu George, Debojyoti Dutta

Thermal Profiling and Modeling of Hadoop Clusters using BigData Applications
Shubbhi Taneja, Ajit Chavan, Yi Zhou, Mohammed Alghamdi, Xiao Qin

11:00-12:30
Session 2

Parallel LDA with Over-Decomposition
Gordon E. Moon, Aravind Sukumaran Rajam, P Sadayappan

11:30-12:30
Keynote Address:
Smart Intersection Control Algorithms for Automated and Connected Vehicles
Sanjay Ranka, University of Florida, Gainesville

14:00-16:00
Session 3

Characterization of Vertex-centric Breadth First Search for Lattice
Ravikant Dindokar, Yogesh Simmhan

BOF: Exploring Opportunities for Big Data Programming Challenge at BigDF

HiPC 2017 WORKSHOP 2:
Computational Fluid Dynamics (CFD)

08:30
Opening of CFD Workshop

08:45-10:00

Keynote Address:
A perspective on the future of CFD and analysis
Aniruddha Mukhopadhyay, PhD, Director, Global Technical Services, ANSYS Inc.

Technical Paper 1 Presentation:
Numerical Simulation of Aerospace Applications using Overset Mesh
Alok Khaware, Abhilash Rajan, Vinay Kumar Gupta

11:00-13:00

Technical Paper 2 Presentation:
Performance Optimization of OpenFOAM on Clusters of Intel® Xeon Phi™ Processors
Ravi Ojha, Prasad Pawar, Sonia Rani, Michael Klemm, Manoj Nambiar

Invited Talk 1: High-resolution modeling of multiscale atmospheric convection
Speaker: Dr. Amit P. Kesarkar, Scientist and Head of Weather and Climate Research Group, National Atmospheric Research Laboratory, India

Invited Talk 2: Simulation driven development of Automotive Lubrication system
Speaker: Dr. Dipak Maiti, President of Indian operations – Simerics Inc.

14:00-15:30

Invited Talk 3: CFD solutions to Engineering Challenges at Shell
Speaker: Dr. Lourens Post, Global Manager of the Fluid Flow & Reactor Engineering team in the Shell R&D

Invited Talk 4: High Performance Computing in CFD
Speaker: Dr. Sudipto Mukhopadhyay, Assistant Professor, Dept. of Mechanical Engineering, IIT Jodhpur

15:30-16:15
Panel Discussion & Closing Remarks & Break

11:00 AM – 1:00 PM

Academic BOF 1: gem5 Development using the POWER ISA v3.0

Speakers:
Basavaraj Talawar, Computer Science and Engineering, NITK, Surathkal.
Sandipan Das and Gautham Shenoy, Linux Technology Center, IBM Systems Development Lab.

Brief Description: The BoF will present the audience with a comprehensive introduction to computer architecture simulation. To improve simulator understanding, we will present the developer view of the state-of-the-art full system simulator, gem5. The participants will be shown how to add new, unimplemented instructions into the POWER processor modules in GEM5 directory structure. Armed with this knowledge, researchers can make modifications at the micro-architectural level in gem5's processor modules.
Schedule:

11:00–11:30: Introduction to Computer Architecture Simulation.

11:30–12:00: gem5 – architecture, modules, code tree.

12:00–1:00:
POWER architecture modules code walk through,
POWER modules implementation,
addition of 2 instructions into the gem5 POWER tree, demo.

Tutorial 1: Introduction to Power Optimization Techniques in HPC

2:00 PM – 4:00 PM

Presenters:
Soham Ghosh, C-DAC, India
Sharda Dixit, C-DAC, India
Ritu Arora, TACC, USA

Tutorial 2: Programming FPGAs for HPC applications using OpenCL

4:30 PM – 6:30 PM

Presenters:
Nitya Hariharan & Kai-Feng Chou, Intel, Corporation

Student Research Symposium (SRS)

Starts at 6:30 PM – Monday, December 18th

The following papers have been accepted for poster presentation on both Day 2 and Day 3 of the conference. They will be on display near the conference industry exhibits, and available for viewing throughout both days. Student authors will be available during breaks to answer questions.

DCTCP Testing: Performance Evaluation of Congestion Control Protocol without ECN
Parth Shukla (CHARUSAT, India) and Bobby Rathore (CHARUSAT, India)

Sherlocker: An Empirical Analysis of Predictive Prefetching for Hybrid Cloud
Sudeep Sharma (IIT Patna, India), Atish Kathpal (Netapp, India), Sourav Basu (Netapp, India) and Jimson Mathew (IIT Patna, India)

GREENER: A Tool for Improving Energy Efficiency of GPU Register File
Vishwesh Jatala (IIT Kanpur, India), Jayvant Anantpur (Mentor Graphics India Pvt Ltd, India) and Amey Karkare (IIT Kanpur, India)

Understanding Uber trip data with Machine-Learning: Geographical Clustering and Time series predictions
A.M.D. Srinivas (M S Ramaiah Institute Of Technology, India), Gagan K Shetty (M S Ramaiah Institute Of Technology, India) and Srinidhi H (M S Ramaiah Institute Of Technology, India)

Massively parallel 3D hemodynamic simulations based on 2D angiogram data to study atherosclerosis in coronary arteries
Madhurima Vardhan (Duke University, USA), John Gounley (Duke University, USA), Amanda Randles (Duke University, USA), James Chen (University of Colorado, USA), Andrew Kahn (UCSD, USA) and Jane Leopold (Harvard University, USA)

Deep Learning for Extracting Noisy Time-Series Signals---Enabling Real-time Multimessenger Astrophysics
Daniel George (National Center for Supercomputing Applications (NCSA), USA), and Eliu Huerta (University of Illinois at Urbana-Champaign, USA)

Persistent Memory Programming Abstractions in Context of Concurrent Applications
Ajay Singh (Indian Institute of Technology Hyderabad, India), Marc Shapiro (INRIA Regal Lab., France) and Gael Thomas (Telecom SudParis, France)

When Polyhedral Optimizations Meet Deep Learning Kernels
Hrishikesh Vaidya (IIT Hyderabad, India), Akilesh B (IIT Hyderabad, India), Abhishek Patwardhan (IIT Hyderabad, India) and Ramakrishna Upadrasta (IIT Hyderabad, India)

Parallel Computing for Iterative Hill Climbing Algorithm to solve TSP
Pramod Yelmewad (National Institute of Technology Karnataka, Surathkal, India), Param Hanji (National Institute of Technology Karnataka, Surathkal, India), Amogha Udupa (National Institute of Technology Karnataka, Surathkal, India), Parth Shah (National Institute of Technology Karnataka, Surathkal, India) and Basavaraj Talawar (National Institute of Technology Karnataka, Surathkal, India)

MBMIR: An efficient Image Reconstruction technique for rapid algorithmic implementation.
Murali Ravi (Sri Sathya Sai Institute of Higher Learning, India), Angu Sewa (Sri Sathya Sai Institute of Higher Learning, India), Shashidhara T G (Analog Devices Inc., India), Sivaramakrishnan S (Analog Devices Inc., India) and Siva Sankara Sai S (Sri Sathya Sai Institute of Higher Learning, India)

Analysis on Traffic Big Data Prediction and Visualization using Parallel M5P and Random Forest Regression
Akhilesh Prasad Mehta (National Institute of Science and Technology, India), T Kiran Kumar (National Institute of Science and Technology, India) and Motahar Reza (National Institute of Science and Technology, India)

Scalable Real Time Rule Engine Model
P Murali Krishna (Sri Sathya Sai Institute of Higher Learning, India) and Pallav Kumar Baruah (Sri Sathya Sai Institute of Higher Learning, India)

Parallelization of the PIC-MCC algorithm for Plasma Simulation on Intel Multicore and Manycore Architectures
Anusha Phadnis (Dhirubhai Ambani Institute of Information and Communication Technology, Gandhinagar, India), Keval Shah (Dhirubhai Ambani Institute of Information and Communication Technology, Gandhinagar, India), Miral Shah (Dhirubhai Ambani Institute of Information and Communication Technology, Gandhinagar, India)and Bhaskar Chaudhury (Dhirubhai Ambani Institute of Information and Communication Technology, Gandhinagar, India)

A Concurrent Graph Object
Sathya Peri (Indian Institute of Technology Hyderabad, India), Muktikanta Sa (Indian Institute of Technology Hyderabad, India)and Nandini Singhal (Indian Institute of Technology Hyderabad, India)

Parallel Implementation Of Screen Content Video Compression based on Block Classification Scheme
Manoj Kumar (Sri Sathya Sai Institute of Higher Learning, India), Digendra Rai (Sri Sathya Sai Institute of Higher Learning, India), Abhishek Gupta (Sri Sathya Sai Institute of Higher Learning, India) and Pallav Kumar Baruah (Sri Sathya Sai Institute of Higher Learning, India)

Improved I/O in DNS code for Cloud Dynamics on 2D processor topology
Lois Thomas (Indian Institute of Tropical Meteorology, India), Bipin Kumar (Indian Institute of Tropical Meteorology, India), Sandeep Jayakumar (Indian Institute of Tropical Meteorology, India), Neethi Suresh (Indian Institute of Tropical Meteorology, India), Suryachandra A Rao (Indian Institute of Tropical Meteorology, India)and Ravi S Nanjundiah (Indian Institute of Tropical Meteorology, India)

Schedule Space Exploration of Nested Parallel Loops
Vandana Kulkarni (PES University, India) and Shilpa Babalad (Indian Institute of Science, India)

A Fast GPU-based Global Sampling Method for Alpha Matting on High Resolution Images
Abhinay B (Sri Sathya Sai Institute of Higher Learning, India)., Anirudh Kalwa (Sri Sathya Sai Institute of Higher Learning, India)., Bhanu Teja Jakkana (Sri Sathya Sai Institute of Higher Learning, India). and Pallav Kumar Baruah (Sri Sathya Sai Institute of Higher Learning, India)

Hardware Acceleration of Boyer Moore Algorithm
Pranav H S (R V College of Engineering, Bangalore, India) and Madhura Purnaprajna (Amrita School of Engineering, Bangalore, India)

Architectural Support for Systems with Non-Volatile Memory
Arpit Joshi (The University of Edinburgh, United Kingdom) Vijay Nagarajan (The University of Edinburgh, United Kingdom), Stratis Viglas (The University of Edinburgh, United Kingdom) and Marcelo Cintra (Intel, Germany)

HBasechianDB2.0 – A Scalable Blockchain Framework on Hadoop Ecosystem
Manuj Subhankar Sahoo (Sri Sathya Sai Institute of Higher Learning, India), Digendra Rai (Sri Sathya Sai Institute of Higher Learning, India), Sai Lakshman P (Sri Sathya Sai Institute of Higher Learning, India), Pallav Kumar Baruah (Sri Sathya Sai Institute of Higher Learning, India)and Adarsh Saraf (Sri Sathya Sai Institute of Higher Learning, India)

Parallel Implementation of Neighbourhood Repulsed Metric Learning
Sathyanarayanan Kumar (Sri Sathya Sai Institute of Higher Learning, India), Chetan Kota (Sri Sathya Sai Institute of Higher Learning, India) and Sai Varun Rachakonda (Sri Sathya Sai Institute of Higher Learning, India)

Optimization of Spark Shuffle Phase
Geetha J (Ramaiah Institute of Technology, India), Chayanika Bhandary (Ramaiah Institute of Technology, India) and Aasia Afreen (Ramaiah Institute of Technology, India)

Accelerating Influence Maximization using Heterogeneous Algorithms
Mridul Haque (Indian Institute of Information Technology Guwahati, India) and Dip Sankar Banerjee (Indian Institute of Information Technology Guwahati)

Partitioning in Apache Spark
J Geetha Reddy (Ramaiah Institute of Technology, India) and H.S. Sreeyuktha (Ramaiah Institute of Technology, India)

GPUScheduler : User Level Preemptive Scheduling for NVIDIA GPUs
Shaleen Garg (International Institute of Information Technology Hyderabad, India), Kishore Kothapalli (International Institute of Information Technology Hyderabad, India), and Suresh Purini (International Institute of Information Technology Hyderabad, India)

Optimizing Molecular Dynamics Acceleration on Parallel Computers
Anu George (Amrita University, India), Noujaz V (Acceleron Labs Pvt. Ltd, India), Narayanan Ms (Acceleron Labs Pvt. Ltd, India), Prasad Lk (Acceleron Labs Pvt. Ltd, India), Prashanth Athri (Amrita University, India), and Madhura Purnaprajna (Amrita University, Bangalore, India)

Deep Learning Strategies for Predicting Iterative Stencil Computations
Aditya Natarajan (College of Engineering Guindy, Anna University, India), Thanasekhar Balaiah (Madras Institute of Technology, Anna University, India), and Ranjani Parthasarathi (College of Engineering Guindy, Anna University, India)

Microbenchmarking Tesla K20 GPU for estimating execution time of a CUDA program
Gargi Alavani (BITS Pilani K. K. Birla Goa Campus, India) and Bharat Arora (BITS Pilani K. K. Birla Goa Campus, India)

Incremental PageRank and k-Core for Dynamic Graphs
Siddharth Jaiswal (Department of Computational and Data Sciences, IISc Bangalore, India), Tilak Naik (Computer Science and Engineering, MS Ramaiah Institute of Technology, India) and Yogesh Simmhan (Department of Computational and Data Sciences, IISc Bangalore, India)

Distributed Querying over Compressed Property Graphs
Swapnil Gandhi (Indian Institute of Science Bangalore, India), Sayandip Sarkar (Indian Institute of Science Bangalore, India), Abhilash Sharma (Indian Institute of Science Bangalore, India) and Yogesh Simmhan (Indian Institute of Science Bangalore, India)

High Performance Algorithms for Learning Exact Bayesian Networks
Subhadeep Karan (University at Buffalo, United States) and Jarosalw Zola (University at Buffalo, United States)

Performance Analysis of Mesh-based NoC's on Routing Algorithms
Amit N Subrahmanya, Allbright D'Souza, Anala M R (RV College of Engineering) , and John Jose (IIT Guwahati)

6:30 PM – 8:00 PM Monday, December 18

Plenary Industry Gala Opening: Exhibits & SRS Posters

Starts at 6:30 PM Monday, December 18

10:00 AM – 5:30 PM Tuesday, December 19 and Wednesday, December 20

Industry/Research Exhibition

Below are the companies who will be providing exhibits and demonstrations at HiPC 2017 Attendees are invited to visit booths and check posted schedules for demonstrations and talks and collect materials these exhibitors have prepared for the conference.

AMD

Google

Hewlett Packard Enterprise

Intel

Microsoft

NetApp

Open Computing

SanDisk

Shell

Xilinx

Day 2 - Tuesday, December 19

Breakfast @ 7:00, Breaks @ 9:30 and 3:00 and 5:15, Lunch @ 12:00 (1 hour)

8:00 AM – 8:30 AM Tuesday, December 19th

Plenary Session: Conference Inauguration and Opening Remarks

8:30 AM – 9:30 AM

HiPC 2017 Keynote Presentation 1

End of Moore's Law: Or, a computer architect's mid-life crisis?

Partha Ranganathan
Google

Abstract: The computer architecture community faces an important and exciting challenge. On one hand, Moore's law is slowing down, stressing traditional assumptions around cheaper and faster systems every year. But at the same time, our demand continues to grow at phenomenal rates, with deeper analysis over ever growing volumes of data, new diverse workloads in the cloud, and smarter edge devices. Is the situation dire, or is this just another phase in the evolution of system architecture. In this talk, I will summarize recent technology and workload trends and discuss likely scenarios for the evolution of future system architectures, with a focus on large "warehouse-scale computing" environments. Specifically, I will discuss how responding to the current opportunities will require more "out-of-the-box" designs —considering the entire datacenter as a computer, co-designing across hardware and software, and developing new architectural constructs.

10:00 AM – 12:00 PM Tuesday, December 19

Technical Session 1: Graph Algorithms

Exact and Parallel Triangle Counting in Dynamic Graphs
Devavret Makkar, David A. Bader, and Oded Green (Georgia Institute of Technology)

Shared-memory Graph Truss Decomposition
Humayun Kabir and Kamesh Madduri (The Pennsylvania State University)

Approximate Computing Techniques for Iterative Graph Algorithms
Ajay Panyala, Omer Subasi, and Mahantesh Halappanavar (Pacific Northwest National Laboratory), Ananth Kalyanaraman (Washington State University), Daniel Chavarría-Miranda (Trovares, Inc.), Sriram Krishnamoorthy (Pacific Northwest National Lab)

Scalable Exact Parent Sets Identification in Bayesian Networks Learning with Apache Spark
Subhadeep Karan and Jaroslaw Zola (University at Buffalo)

Parallel Exact Dynamic Bayesian Network Structure Learning with Application to Gene Networks
Vasimuddin Md (Indian Institute of Technology Bombay), Srinivas Aluru (Georgia Institute of Technology)

Parallel Asynchronous Distributed-Memory Maximal Independent Set Algorithm with Work Ordering
Thejaka Amila Kanewala and Marcin Zalewski (Indiana University), Andrew Lumsdaine (Pacific Northwest Natl Lab and U Washington)

10:00 AM – 12:00 PM Tuesday, December 19

IRUS 1: Practical Applications of Machine Learning in the e-commerce sector

Anil R. Yelundur, Applied Scientist with Amazon
Presentation # 1 : Anomaly Detection using Tensor Decomposition

Samik Datta, Principal Data Scientist, Flipkart, India
Presentation # 2 : Personalization via Persona-ization for Merchandising Lifestyle Articles

Ashish Kulkarni, Machine Learning Scientist, Amazon, India
Presentation # 3 : Answering Questions on Products

Saurabh, Machine Learning Scientist, Amazon, India
Presentation # 4 : Estimating the performance of online A/B tests on historical logs

10:00 AM – 12:00 PM Tuesday, December 19

Sponsored Tech Session: SanDisk

Title: Accelerating your workloads with NVMe SSDs

Join Esther Spanjer from Western Digital for this presentation where she will discuss how NVMe SSDs are enabling the next generation datacenter. She will discuss the unique advantages of NVMe SSDs over alternative solutions and discuss the best use cases for NVMe SSD in a variety of environments, such as vSAN, Hadoop, Microsoft Storage Spaces Direct, SQL Server and others. She will furthermore discuss the next development in the datacenter when NVMe Over Fabric will be introduced, and how this can enable new architectures that allow for higher performance at lower cost.

1:00 PM – 3:00 PM Tuesday, December 19

Technical Session 2: Architecture and Communication

Designing Registration Caching Free High-Performance MPI Library with Implicit On-Demand Paging (ODP) of InfiniBand
Mingzhe Li, Xiaoyi Lu, Hari Subramoni, and Dhabaleswar Panda (The Ohio State University)

Last Level Collective Hardware Prefetching For Data-Parallel Applications
George Michelogiannakis and John Shalf (Lawrence Berkeley National Laboratory)

Kernel-assisted Communication Engine for MPI on Emerging Manycore Processors
Jahanzeb Maqbool Hashmi, Khaled Hamidouche, Hari Subramoni, and Dhabaleswar Panda (The Ohio State University)

Support for Power Efficient Proactive Cooling Mechanisms
Bilge Acun (University of Iliinois at Urbana-Champaign), Eun Kyung Lee (IBM T.J. Watson Research Center), Yoonho Park (IBM T.J. Watson Research Center), Laxmikant Kale (University of Illinois at Urbana-Champaign)

Redundant Arithmetic based High Speed Carry Free Hybrid Adders with Built-In Scan Chain on FPGAs
Ayan Palchaudhuri and Anindya Sundar Dhar (Indian Institute of Technology Kharagpur)

ConvLight: A Convolutional Neural Network Accelerator with Neuromorphic Photonic-based Computing
Dharanidhar Dang, Jyotikrishna Dass, and Rabi Mahapatra (Texas A&M University)

1:00 PM – 3:00 PM Tuesday, December 19

Sponsored Tech Session: Microsoft

Title: Microsoft Azure and HPC

Speakers: Tony Wu, Rangarajan Srirangam, and Rakesh Patil

Azure provides on-demand true HPC capability that enables you to run large parallel and batch HPC jobs and extend your on-premises HPC cluster to the cloud when you need more capacity. In this session we will start by introducing the Azure cloud and the Azure presence in India. After this we will do a deep dive of HPC in Azure cloud including Hardware (compute, storage, and network) and software (Middleware and ISV application) stacks along with typical solution architectures in both hybrid and native cloud deployments. We will explore HPC solutions in practice by considering use case scenarios in key industries including finance, manufacturing, life sciences, and educational research.

3:15 PM – 5:15 PM Tuesday, December 19

Technical Session 3: Algorithms

Provably Efficient Scheduling of Dynamically Allocating Programs on Parallel Cache Hierarchies
Harsha Vardhan Simhadri (Microsoft Research), Guy Blelloch and Phillip Gibbons (Carnegie Mellon University)

Further Explorations in State-Space Search for Optimal Task Scheduling
Michael Orr and Oliver Sinnen (University of Auckland)

A Novel Approach for Job Scheduling Optimizations under Power Cap for ARM and Intel HPC Systems
Dineshkumar Rajagopal (Bull Atos Technologies), Daniele Tafani (Leibniz Supercomputing Centre), Yiannis Georgiou and David Glesser (Bull Atos Technologies), Michael Ott (Leibniz Supercomputing Centre)

A Memory Congestion-aware MPI Process Placement for Modern NUMA Systems
Mulya Agung, Muhammad Alfian Amrizal, Kazuhiko Komatsu, Ryusuke Egawa, and Hiroyuki Takizawa (Tohoku University)

Expander: Lock-free Cache for a Concurrent Data Structure
Pooja Aggarwal (IBM Research, India), Smruti Sarangi (IIT Delhi)

Adaptive Code Refinement: A Compiler Technique and Extensions to Generate Self-Tuning Applications
Maxime Schmitt and Philippe Helluy (Univ. of Strasbourg, INRIA), Cédric Bastoul (University of Strasbourg)

3:15 PM – 5:15 PM Tuesday, December 19

IRUS 2: Performance Challenge in Deep Learning and Scientific Computing

Ashish Sirasao
Distinguished Engineer, Xilinx Software and IP Team, USA
Presentation # 1: Deep Learning using Xilinx FPGAs

Servesh Muralidharan
Senior Fellow, CERN IT Department.
Presentation # 2 : Computing Challenges in HEP for WLHC grid

5:30 PM – 7:30 PM Tuesday, December 19

Plenary Panel: HPC in India

As a special feature of HiPC 2017, a session on HPC initiatives in India is being organized and will feature presentations by leading researchers from academia, industry and Indian government R&D labs.

R. Govindarajan **Indian Institute of Science, Bangalore**	*Moderator*
Raghunathan Ramakrishnan **Tata Institute of Fundamental Research, Hyderabad**	*Accelerating scientific discoveries with Computational Chemistry Big Data*
N. Balakrishnan **Indian Institute of Science, Bangalore**	*Indian CFD at crossroads*
Ravi S. Nanjundiah **Indian Institute of Tropical Meteorology, Pune**	*To Be Announced*
Director General, C-DAC Representative	*Supercomputing in C-DAC*

Day 3 – Wednesday, December 20

Breakfast @ 7:30, Breaks @ 9:30 and 3:00, Lunch @ 12:00 (1 hour)

8:30 AM – 9:30 AM Wednesday, December 20

HiPC 2017 Keynote Presentation 2

Machine Learning @ Amazon

Rajeev Rastogi

Amazon India

Abstract: In this talk, I will first provide an overview of key problem areas where we are applying Machine Learning (ML) techniques within Amazon such as product demand forecasting, product search, and information extraction from reviews, and associated technical challenges. I will then talk about three specific applications where we use a variety of methods to learn semantically rich representations of data: question answering where we use deep learning techniques, product size recommendations where we use probabilistic models, and fake reviews detection where we use tensor factorization algorithms. I will point out the computing challenges associated with these applications and how parallelism can be exploited to scale to large datasets.

10:00 AM – 12:00 PM Wednesday, December 20

Technical Session 4: Big Data, Machine Learning and Optimization

Parallel Deep Convolutional Neural Network Training by Exploiting the Overlapping of Computation and Communication
Sunwoo Lee, Dipendra Jha, Ankit Agrawal, Alok Choudhary, and Wei-Keng Liao (Northwestern University)

Parallel Dynamic Data Driven Approaches for Synthetic Aperture Radar
Adeesha Wijayasiri, Tania Banerjee, Sanjay Ranka, Sartaj Sahni, and Mark Schmalz (University of Florida – Gainesville Campus)

ARM Wrestling with Big Data: A Study of Commodity ARM Server for Big Data Workloads
Jayanth Kalyanasundaram and Yogesh Simmhan (Indian Institute of Science)

MPI-LiFE: Designing High-Performance Linear Fascicle Evaluation of Brain Connectome with MPI
Shashank Gugnani and Xiaoyi Lu (The Ohio State University), Franco Pestilli and Cesar Caiafa (Indiana University), Dhabaleswar Panda (The Ohio State University)

Reducing network congestion and global communication bottlenecks during aggregation on Torus and Dragonfly topologies for writing hierarchical data
Sidharth Kumar and Duong Hoang (SCI, University of Utah), Steve Petruzza (SCI, University of Utah and University of Rome "Tor Vergata"), John Edwards (Idaho State University), Valerio Pascucci (SCI, University of Utah)

Fast Parallel Randomized QR with Column Pivoting Algorithms for Reliable Low-rank Matrix Approximations

Jianwei Xiao and Ming Gu (UC Berkeley), Julien Langou (University of Colorado Denver

1:00 PM – 3:00 PM Wednesday, December 20

Technical Session 5: Graph Algorithms and GPU

An X10 based Distributed Streaming Graph Database Engine
Miyuru Dayarathna (WSO2 Inc.), Sathya Bandara, Nandula Jayamaha, Mahen Herath, Achala Madhushan, and Sanath Jayasena (University of Moratuwa), Toyotaro Suzumura (IBM T.J. Watson Research Center)

GPU-centric Communication on NVIDIA GPU Clusters with InfiniBand: A Case Study with OpenSHMEM
Sreeram Potluri, Anshuman Goswami, and Davide Rossetti (NVIDIA Corporation), Manjunath Gorentla Venkata and Neena Inam (Oak Ridge National Laboratory), Chris J. Newburn (NVIDIA Corporation)

Distributed Algorithm for High-Utility Subgraph Pattern Mining Over Big Data Platform
Alind Khare, Vikram Goyal, and Srikant Baride (IIIT-Delhi), Michael McDermott, Dhara Shah, and Sushil Prasad (Georgia State University)

ReCALL: Reordered Cache aware LocaLity based Graph Processing
Kartik Lakhotia, Shreyas Singapura, Rajgopal Kannan, and Viktor Prasanna (University of Southern California)

Characterization of data movement requirements for sparse matrix computations on GPUs
Sureyya Emre Kurt, Vineeth Thumma, Changwan Hong, Aravind Sukumaran-Rajam, and Sadayappan P (The Ohio State University)

1:00 PM – 3:00 PM Wednesday, December 20

Sponsored Tech Session: Xilinx

Title : Hardware Accelerators for NoSQL databases
Speaker: Chidamber Kulkarni – RENIAC

In this talk, we will discuss our experiences in building a low-latency, high-throughput data acceleration engine for NoSQL databases, such as Cassandra. We will introduce the concept of a Data Proxy that marries the concepts of a transparent data layer proxy, that can talk CQL (for Cassandra), and the concept of a storage engine, that implements caching of data across multiple memory technologies such as SRAM, DRAM, and Flash memories. We will then discuss the different accelerator technologies that can be used to build such a data acceleration engine and elucidate why FPGAs make the right choice. In this talk, we will also highlight the key design decisions and trade-offs that we made to leverage the power of FPGA devices to deliver 1/3rd to 1/10th lower latency and significantly higher throughput by saturating the 10Gb Ethernet within a server. Finally, we will present some of our benchmark results to validate our claims. To conclude, we will show how the Data Proxy can be applied to other NoSQL databases, and adjacent technologies such as Search (using Solr/ElasticSearch) both as an on-prem solution and on cloud platforms with FPGAs, such as AWS F1.

Title : Video codec acceleration on FPGAs
Speaker: Mahesh N Shukla – Xilinx

With the need to handle high-resolution video streams (1080p to 4K and beyond), advanced video compression standards such as HEVC, VP9, and AV1 are being developed and deployed. The high resolution raw streams present a huge challenge, with the need to compress enormous data in the available network bandwidth. For example, 4kp60-10bit takes around 5Gbps in its raw form, which means that in a network bandwidth of 10Mbps, raw data needs to be compressed by 500x. To cater to this need, the newer compression standards are more complex and are challenging to implement. For example, there is a 30x increase in complexity from H264 1080p30 to HEVC 4kp60. This complexity means that hardware / SoC implementations need higher area and power budgets. FPGAs are well suited to address this problem, as they are programmable. In this talk, we will discuss the video acceleration strategy at Xilinx and the enablement of solutions for video applications. We will also talk about the video codec roadmap and how dynamic codec switching offers high performance at a significantly low cost.

3:30 PM – 5:30 PM Wednesday, December 20

Plenary Panel: Machine Learning and Mathematical Modeling – Towards Robust, Reliable and Certifiable Machine Learning

This panel will draw parallels to the analysis of traditional mathematical models and illustrate potential research pathways in machine learning. We will use a case study approach with distinguished panelists from traditional machine deep learning and mathematical modeling domains to explore the complementarity attributes of these domains, and the opportunities (and challenges) in bringing together these currently disparate communities and ideas.

Abani Patra
Professor and Director of Computational, Data Sciences and Engineering,
University at Buffalo
Manish Parashar
Distinguished Professor of Computer Science and Director of Rutgers Discover Institute,
Rutgers University
Valerio Pascucci
Professor of Computer Science and Director of the Center for Extreme Data Management Analysis and Visualization, University of Utah
Dimitri Kusnezov
Chief Scientist & Senior Advisor to the Secretary of the U.S. Department of Energy
Ashish Kapoor
Microsoft Research

Starting at 6:30 PM Wednesday, December 20

Banquet & Awards

Details to be announced

Day 4 – Thursday, December 21

Breakfast @ 7:30, Break@ 9:30, Lunch @ 12:00

8:30 AM – 9:30 AM Thursday, December 21

HiPC 2017 Keynote Presentation 3

Computing Just What You Need: Online Data Analysis and Reduction at Extreme Scales

Ian Foster
Argonne National Lab and University of Chicago

Abstract: A growing disparity between supercomputer computation speeds and I/O rates makes it increasingly infeasible for applications to save all results for offline analysis. Instead, applications must analyze and reduce data online so as to output only those results needed to answer target scientific question(s). This change in focus complicates application and experiment design and introduces algorithmic, implementation, and programming model challenges that are unfamiliar to many scientists and that have major implications for the design of various elements of supercomputer systems. I review these challenges and describe methods and tools that various groups, including mine, are developing to enable experimental exploration of algorithmic, software, and system design alternatives.

10:00 AM – 12:00 PM Thursday, December 21

Technical Session 6: System Software

Integrating External Resources with a Task-Based Programming Model
Zhihao Jia (Stanford University), Sean Treichler (NVIDIA Research), Galen Shipman (Los Alamos National Laboratory), Michael Bauer (NVIDIA Research), Noah Watkins and Carlos Maltzahn (UC Santa Cruz), Pat McCormick (Los Alamos National Laboratory), Alex Aiken (Stanford University)

Enabling Dependability-Driven Resource Use and Message Log-Analysis for Cluster System Diagnosis
Edward Chuah (The Alan Turing Institute & The University of Warwick), Arshad Jhumka (University of Warwick), Samantha Alt (Intel Corporation), Theo Damoulas (The Alan Turing Institute & The University of Warwick), Nentawe Gurumdimma (The University of Jos), Marie-Christine Sawley (Intel Corporation), Bill Barth (University of Texas at Austin), Tommy Minyard (The Texas Advanced Computing Center), James Browne (University of Texas at Austin)

Context-Aware Memory Profiling for Speculative Parallelism
Changsu Kim, Juhyun Kim, and Juwon Kang (POSTECH), Jae W. Lee (Seoul National University), Hanjun Kim (POSTECH)

Lifting Barriers Using Parallel Polyhedral Regions
Harenome Ranaivoarivony-Razanajato, Cédric Bastoul, and Vincent Loechner (CAMUS team, INRIA Nancy Grand-Est and University of Strasbourg)

Exploiting Common Neighborhoods to Optimize MPI Neighborhood Collectives
Seyed Hessamedin Mirsadeghi (Queen's University), Jesper Larsson Traff (Vienna University of Technology/TU Wien), Pavan Balaji (Argonne National Laboratory), Ahmad Afsahi (Queen's University)

Efficient Fork-Join on GPUs through Warp Specialization
Arpith Jacob, Alexandre Eichenberger, and Hyojin Sung (IBM T.J. Watson Research Center), Samuel Antao (IBM Research UK), Gheorghe-Teodor Bercea and Carlo Bertolli (IBM T.J. Watson Research Center), Alexey Bataev, Tian Jin, and Tong Chen (IBM Research), Zehra Sura, Rokos Georgios, and Kevin O'Brien (IBM T.J. Watson Research Center)

10:00 AM – 12:00 PM Thursday, December 21

Academic BOF 2: Parallel LSG Computation

Organizers:
John Augustine, IIT Madras & Kishore Kothapalli, IIIT Hyderabad

(Check website for more details.)

1:00 PM – 3:00 PM Thursday, December 21

Technical Session 7: GPU Frameworks and Applications

Thrust++: Extending Thrust Framework for Better Abstraction and Performance
Ajai George, Sankar Manoj, and Sanket Gupte (BITS Pilani K K Birla Goa Campus), Sayantan Mitra (Siemens Technology & Services Pvt Ltd, Bangalore, India), Santonu Sarkar (BITS Pilani K K Birla Goa Campus)

A Novel Implementation of 2D3V Particle-In-Cell (PIC) Algorithm for Kepler GPU Architectures
Harshil Shah, Siddharth Kamaria, Riddhesh Markandeya, Miral Shah, and Bhaskar Chaudhury (DA-IICT, Gandhinagar)

Parallelizing Hines Matrix Solver in Neuron Simulations on GPU
Dharma Teja Vooturi and Kishore Kothapalli (International Institute of Information Technology, Hyderabad), Upinder Bhalla (National Centre for Biological Sciences)

Building Halo Merger Trees from the Q Continuum Simulation
Esteban Rangel (Northwestern University), Nicholas Frontiere (University of Chicago), Salman Habib and Katrin Heitmann (Argonne National Laboratory), Wei-Keng Liao (Northwestern University), Ankit Agrawal (Iowa State University), Alok Choudhary (Northwestern University)

A Memory-Efficient GPU Method for Hamming and Levenshtein Distance Similarity
Andrew Todd (University of Missouri), Marziyeh Nourian and Michela Becchi (North Carolina State University)

2018

December 17-20, 2018

Bengaluru, India