



2016 International Conference on Parallel Processing

■ August 16 – 19, 2016 ■ Temple University Center City ■ Philadelphia, PA, USA

CONFERENCE AT A GLANCE

	TUESDAY AUG. 16	WEDNESDAY AUG. 17		THURSDAY AUG. 18		FRIDAY AUG. 19	
7:30	Registration (Room 603)	Registration (Room 220)		Registration (Room 220)		Registration (Room 220)	
8:00	Workshops	Opening Remarks					
8:30		Keynote I Speaker: Chita Das		Keynote II Speaker: George Pappas		Keynote III Speaker: Ricardo Bianchini	
9:00		Refreshments		Refreshments		Refreshments	
9:30	Refreshments	Refreshments		Refreshments		Refreshments	
10:00	Workshops	Session 1A	Session 1B	Session 4A	Session 4A	Session 7A	Session 7B
10:30		Lunch		Lunch		Lunch	
11:00							
11:30							
12:00	Lunch	Lunch		Lunch		Lunch	
12:30							
1:00							
1:30	Workshops	Session 2A	Session 2B	Session 5A	Session 5B	Session 8A	Session 8B
2:00		Refreshments		Refreshments		Recess	
2:30							
3:00							
3:30	Refreshments	Refreshments		Refreshments		Recess	
4:00	Workshops	Session 3A	Session 3B	Session 6A	Session 6B		
4:30		Recess		Recess			
5:00							
5:30	Recess	ICPP Reception Estia Restaurant 6:00 – 8:00pm		ICPP Banquet Joy Tsin Lau Restaurant 6:30 – 9:00pm			
6:00							
6:30							
7:00							
7:30							
8:00	ICPP Reception Estia Restaurant 6:00 – 8:00pm		ICPP Banquet Joy Tsin Lau Restaurant 6:30 – 9:00pm				
6:30							
7:00							
7:30							
8:00							

45th Annual International Conference on Parallel Processing

PROGRAM

Tuesday, Aug. 17

Opening Remarks

Room 222: 8:00 – 8:30

Keynote I

Room 222: 9:30 – 10:30

Chair – Jie Wu

“ICPP: Then, Now and the Future”

Chita Das

9:30 – 10:00 Break

Session 1A: Data & Cloud I

Room 408: 10:00 – 12:00

Chair – Michael Scott

- Efficient Virtual Network Embedding for Variable Size Virtual Machines in Fat-tree Data Centers, Jun Duan, Yuanyuan Yang
- MIC: An Efficient Anonymous Communication System in Data Center Networks, Tingwei Zhu, Fang Wang, Yu Hua, Dan Feng, Qingyu Shi, Jiahao Liu
- AppBag : Application-aware Bandwidth Allocation for Virtual Machines in Cloud Environment, Dian Shen, Junzhou Luo, Fang Dong, Junxue Zhang
- Performance Boosting Opportunities under Communication Imbalance in Power-Constrained HPC Clusters, Leonardo Piga, Indrani Paul, Wei Huang
- RRect: A Novel Structure With High Availability and Power Efficiency for Data Centers, Zhenhua Li, Yuanyuan Yang

Session 1B: Architecture I

Room 421: 10:00 – 12:00

Chair – Dong Li

- Making In-Memory Frequent Pattern Mining Durable and Energy Efficient, Yi Lin, Po-Chun Huang, Duo Liu, Xiao Zhu, Liang Liang
- Run-time Performance Estimation and Fairness-oriented Scheduling Policy for Concurrent GPGPU Applications, Hu Qingda, Shu Jiwu, Fan Jie, Lu Youyou
- Performance Analysis of GPU-based Convolutional Neural Networks, Xiaqing Li, Guangyan Zhang, H. Howie Huang, Zhufan Wang, Weimin Zheng

- Proxy-Guided Load Balancing of Graph Processing Workloads on Heterogeneous Clusters, Shuang Song, Meng Li, Xinnian Zheng, Michael Lebeane, Jee Ho Ryoo, Reena Panda, Andreas Gerstlauer, Lizy K. John
- Piccolo: A Fast and Efficient Rollback System for Virtual Machine Clusters, Lei Cui, Zhiyu Hao, Chonghua Wang

12:00 – 1:30 Lunch Break

Session 2A: Parallel Algorithms I

Room 408: 1:30 – 3:30

Chair – Peng Li

- Parallel k-means++ for Multiple Shared-Memory Architectures, Patrick Mackey, Robert Lewis
- High Performance Parallel Algorithms for the Tucker Decomposition of Sparse Tensors, Oguz Kaya, Bora Ucar
- Parallel Tree Traversal for Nearest Neighbor Query on the GPU, Moohyeon Nam, Jinwoong Kim, Beomseok Nam
- Resilient application co-scheduling with processor redistribution, Anne Benoit, Loic Pottier, Yves Robert
- Efficient Parallel Algorithms for k-Center Clustering, Jessica McClintock, Anthony Wirth

Session 2B: Architecture II

Room 421: 1:30 – 3:30

Chair – Federico Silla

- Understanding the Architectural Characteristics of EDA Algorithms, Xin Wang, Yunping Lu, Yi Li, Weihua Zhang
- Exploring Variation-Aware Fault-Tolerant Cache under Near-Threshold Computing, Jing Wang, Yanjun Liu, Weigong Zhang, Keni Qiu, Xin Fu, Tao Li
- Tetris Write: Exploring More Write Parallelism Considering PCM Asymmetries, Li Zheng, Wang Fang, Feng Dan, Hua Yu, Tong Wei, Liu Jingning, Liu Xiang
- ROP: Alleviating Refresh Overheads via Reviving the Memory System in Frozen Cycles, Ping Huang, Wenjie Liu, Kun Tang, Xubin He, Ke Zhou

- Thread Similarity Matrix: Visualizing Branch Divergence in GPGPU programs, Zhibin Yu, Lieven Eeckhout, Chengzhong Xu

3:30 – 4:00 Break

Session 3A: Programming Techniques I

Room 408: 4:00 – 5:30

Chair – Yves Robert

- One-Sided Interface for Matrix Operations using MPI-3 RMA: A Case Study with Elemental, Sayan Ghosh, Jeff Hammond, Antonio J. Peña, Pavan Balaji, Assefaw Gebremedhin, Barbara Chapman
- SWAP-Assembler 2: Optimization of De Novo Genome Assembler at Extreme Scale, Jintao Meng, Sangmin Seo, Pavan Balaji, Yanjie Wei, Bingqiang Wang, Shenzhong Feng
- Programming Techniques for the Automata Processor, Indranil Roy, Ankit Srivastava, Srinivas Aluru
- RCHC: a Holistic Runtime System for Concurrent Heterogeneous Computing, Jinsu Park, Woongki Baek

Session 3B: Parallel Algorithms II

Room 421: 4:00 – 5:30

Chair – Michael Raynal

- An Unbounded Nonblocking Double-ended Queue, Matthew Graichen, Joseph Izraelevitz, Michael Scott
- Criticality-Aware Partitioning for Multicore Mixed-Criticality Systems, Jian-Jun Han, Xin Tao, Dakai Zhu, Hakan Aydin
- A Parallel Hill-Climbing Refinement Algorithm for Graph Partitioning, Dominique Lasalle, George Karypis
- Massively-Parallel Lossless Data Decompression, Evangelia Sitaridi, Rene Mueller, Tim Kaldewey, Guy Lohman, Kenneth Ross

ICPP '16 Reception

Estia Restaurant

1405 Locust Street

6:00 – 8:00

Thursday, Aug. 18

Keynote II

Room 222: 8:30 – 9:30

Chair – Krishna Kant

“Internet of Things: Challenges and Opportunities”

George Pappas

9:30 – 10:00 Break

Session 4A: Datacenter & Cloud II

Room 408: 10:00 – 12:00

Chair – Taisuke Boku

- Improving Data Transfer Throughput with Direct Search Optimization, Prasanna Balaprakash, Vitali Morozov, Rajkumar Kettimuthu., Kalyan Kumaran, Ian Foster
- MPI Overlap: Benchmark and Analysis, Alexandre Denis, François Trahay
- High Performance MPI Library for Container-based HPC Cloud on InfiniBand Clusters, Jie Zhang, Xiaoyi Lu, Dhableswar Panda
- AccuracyTrader: Accuracy-aware Approximate Processing for Low Tail Latency and High Result Accuracy in Cloud Online Services, Rui Han, Siguang Huang, Fei Tang, Fugui Chang, Jianfeng Zhan
- CoARC: Co-operative, Aggressive Recovery and Caching for Failures in Erasure Coded Hadoop, Pradeep Subedi, Ping Huang, Tong Liu, Joseph Moore, Stan Skelton, Xubin He

Session 4B: Cyberphysical Systems I

Room 421: 10:00 – 12:00

Chair – David Du

- Optimal Multi-Taxi Dispatch for Mobile Taxi-Hailing Systems, Guoju Gao, Mingjun Xiao, Zhenhua Zhao
- Fast RFID Polling Protocols, Jia Liu, Bin Xiao, Xuan Liu, Lijun Chen
- Exploiting Real-time Traffic Light Scheduling with Taxi Traces, Zongjian He, Daqiang Zhang, Jiannong Cao, Xuefeng Liu, Xiaopeng Fan, Chengzhong Xu
- An Efficient Wireless Power Transfer System To Balance the State of Charge of Electric Vehicles, Ankur Sarker, Chenxi Qiu, Haiying Shen, Andrea Gill, Joachim Taiber, Mashrur Chowdhury, Jim Martin, Mac Devine, Aj Rindos

- EchoTrack : Accurate Device-free Hand Tracking on Smart Phones, Huijie Chen, Fan Li, Yu Wang

12:00 – 1:30 Lunch Break

Session 5A: Parallel Algorithms III

Room 408: 1:30 – 3:30

Chair – Jenq-Kuen Lee

- Randomly Optimized Grid Graph for Low-Latency Interconnection Networks, Koji Nakano, Daisuke Takafuji, Satoshi Fujita, Hiroki Matsutani, Ikki Fujiwara, Michihiro Koibuchi
- Optimal Collision/Conflict-free Distance-2 Coloring in Synchronous Broadcast/Receive Tree Networks, Davide Frey, Hicham Lakhlef, Michel Raynal
- Help-optimal and Language-portable Lock-free Concurrent Data Structures, Bapi Chatterjee, Ivan Walulya, Philippas Tsigas
- DC-Top-k: A Novel Top-k Selecting Algorithm and Its Parallelization, Zhengyuan Xue, Ruixuan Li, Heng Zhang, Xiwu Gu, Zhiyong Xu
- Efficient 2-body Statistics Computation on GPUs: Parallelization & Beyond, Napath Pitaksiranan, Zhila Nouri, Yi-Cheng Tu

Session 5B: Storage Systems

Room 421: 1:30 – 3:30

Chair – Xubin He

- Think Global, Act Local: A Buffer Cache Design for Global Ordering and Parallel Processing in the WAFL File System, Peter Denz, Matthew Curtis-Maury, Vinay Devadas
- Improving RAID Performance Using an Endurable SSD Cache, Chu Li, Dan Feng, Yu Hua, Fang Wang
- In situ Storage Layout Optimization for AMR Spatio-temporal Read Accesses, Houjun Tang, Suren Byna, Steve Harenberg, Wenzhao Zhang, Xiaocheng Zou, Daniel F. Martin Bin Dong, Dharshi Devendran Kesheng Wu, David Trebotich, Scott Klasky, Nagiza Samatova
- Managing I/O Interference in a Shared Burst Buffer System, Sagar Thapaliya, Purushotham Bangalore, Jay Lofstead, Kathryn Mohror, Adam Moody

- Guaranteed Bang for the Buck: Modeling VDI Applications with Guaranteed Quality of Service, Hao Wen, David Du, Milan Shetti, Doug Voigt, Shanshan Li

3:30 – 4:00 Break

Session 6A: Programming

Techniques II

Room 408: 4:00 – 5:30

Chair – Masahiro Yasugi

- Scalable Hierarchical Polyhedral Compilation, Benoit Pradelle, Benoit Meister, Muthu Manikandan Baskaran, Athanasios Konstantinidis, Thomas Henretty, Richard Lethin
- The Future(s) of Transactional Memory, Jingna Zeng, Joao Barreto Seif Haridi, Luis Rodrigues, Paolo Romano
- Declarative Tuning for Locality in Parallel Programs, Sanjay Chatterjee, Nick Vrvilo, Zoran Budimlic, Kath Knobe, Vivek Sarkar
- Ensemble Toolkit: Scalable and Flexible Execution of Ensembles of Tasks, Vivekanandan Balasubramanian, Antons Treikalis, Ole Weidner, Shantenu Jha

Session 6B: Cyberphysical Systems II

Room 421: 4:00 – 5:30

Chair – Fang Wang

- TECH: A Thermal-Aware and Cost Efficient Mechanism for Colocation Demand Response, Ziqi Zhao, Fan Wu, Shaolei Ren, Xiaofeng Gao, Guihai Chen, Yong Cui
- A Scalability Comparison Study of Data Management Approaches for Smart Metering Systems, Housseem Chihoub, Christine Collet
- A Comparison of Accelerator Architectures for Radio-Astronomical Signal-Processing Algorithms, John W. Romein
- MobiSensing: Exploiting Human Mobility for Multi-Application Mobile Data Sensing with Low User Intervention, Kang Chen, Haiying Shen

ICPP '16 Banquet

Joy Tsin Lau Restaurant

1026 Race Street

6:30 – 9:00

Friday, Aug. 19

Keynote Address

Room 222: 9:30 – 10:30

Chair – Hong Jiang

“Improving Efficiency in Large-Scale Datacenters: Experience and Lessons”

Ricardo Bianchini

10:30 – 11:00 Break

Session 7A: Performance Modeling

Room 408: 10:00 – 12:00

Chair – Xizhou Feng

- Sparse Matrix Format Selection with Multiclass SVM for SpMV on GPU, Akrem Benatia, Weixing Ji, Yizhuo Wang, Feng Shi
- On the Impact of Widening Vector Registers on Sequence Alignment, Jeff Daily, Ananth Kalyanaraman, Sriram Krishnamoorthy, Bin Ren
- The Case for Cross-Component Power Coordination on Power Bounded Systems, Rong Ge, Xizhou Feng, Yangyang He, Pengfei Zou
- Performance Maximization via Frequency Oscillation on Temperature Constrained Multicore Processors, Shi Sha, Wujie Wen, Ming Fan, Shaolei Ren, Gang Quan
- RMD: A Resemblance and Mergence based Approach for High Performance Deduplication, Panfeng Zhang, Ping Huang, Xubin He, Hua Wang, Lingyu Yan, Ke Zhou

Session 7B: GPU Applications

Room 421: 10:00 – 12:00

Chair – TBA

- GFlink: An In-Memory Computing Architecture on Heterogeneous CPU-GPU Clusters for Big Data, Kenli Li, Cen Chen, Aijia Ouyang, Keqin Li
- Partial Flattening: A Compilation Technique for Irregular Nested Parallelism on GPGPUs, Ming-Hsiang Huang, Wu Yang
- RegTT: Accelerating Tree Traversals on GPUs by Exploiting Regularities, Feng Zhang, Peng Di, Hao Zhou, Jingling Xue
- Optimizing GPU Register Usage: Extensions to OpenACC and Compiler Optimizations, Xiaonan Tian, Dounia Khaldi, Rengan Xu, Barbara Chapman

- HppCnn: A High-Performance, Portable Deep-Learning Library for GPGPUs, Yi Yang, Min Feng, Srimat Chakradhar

12:00 – 1:30 Lunch Break

Session 8A: Applications

Room 408: 1:30 – 2:45

Chair – Anne Benoit

- Locality-aware laplacian mesh smoothing, Guillaume Aupy, Jeonghyung Park, Raghavan Padma
- Fault Tolerant Support Vector Machines, Sameh Shohdy, Abhinav Vishnu, Gagan Agrawal
- Parallel Two-Dimensional Unstructured Anisotropic Delaunay Mesh Generation of Complex Domains for Aerospace Applications, Juliette Pardue, Andrey Chernikov

Session 8B: Scalable Software

Room 421: 1:30 – 2:45

Chair – Ping Huang

- PARVMEC: An Efficient, Scalable Implementation of the Variational Moments Equilibrium Code, Sudip Seal, Steven Hirshman, Andreas Wingen, Robert Wilcox, Mark Cianciosa, Ezekial Unterberg
- RepEx: A Flexible Framework for Scalable Replica Exchange Molecular Dynamics Simulations, Antons Treikalis, Andre Merzky, Haoyuan Chen Tai-Sung Lee, Darrin York, Shantenu Jha
- PCAF: Scalable, High Precision k-NN Search using Principal Component Analysis based Filtering, Huan Feng, David Eysers, Steven Mills, Yongwei Wu, Zhiyi Huang

WORKSHOPS

Tuesday, Aug. 19

Room 601	8:00AM - 5:30PM: #1 - P2S2
Room 617	8:00AM - 5:30PM: #2 - DSE
Room 606	8:00AM - 12:00PM: #3 - HPC4BD 1:00PM - 5:30PM: #4 - AWASN
Room 607	8:00AM - 12:00PM: #5 - HUCAA 1:00PM - 5:30PM: #6 - PSTI
Room 608	8:00AM - 12:00PM: #7 - SHPCloud 1:00PM - 5:30PM: #8 - EMS
Room 610	8:00AM - 12:00PM: #9 - SRMPDS 1:00PM - 5:30PM: #10 - PASA

1. International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2): <http://www.mcs.anl.gov/events/workshops/p2s2/2016>.
2. International Workshop on Data Science and Engineering (DSE): http://www.ieee-ccsna.org/dse_icpp2016/.
3. International Workshop on High Performance Computing for Big Data (HPC4BD): <http://people.sabanciuniv.edu/kaya/hpc4bd/hpc4bd.html>.
4. International Workshop on Applications of Wireless Ad hoc and Sensor Networks (AWASN): <http://www.comp.sd.tmu.ac.jp/kazuyasakai/events/awasn16.html>.
5. International Workshop on Heterogeneous and Unconventional Cluster Architectures and Applications (HUCAA): <http://www.hucaa-workshop.org>.
6. International Workshop on Parallel Software Tools and Tool Infrastructures (PSTI): <http://www.psti-workshop.org/2016>.
7. International Workshop on Sustainable HPC Cloud (SHPCloud): <https://sites.google.com/a/temple.edu/shpc2016>.
8. International Workshop on Embedded Multicore Systems (EMS): <https://sites.google.com/site/icppems2016>.
9. International Workshop on Scheduling and Resource Management for Parallel and Distributed Systems (SRMPDS): <https://sites.google.com/site/srmpds>.
10. International Workshop on Power-aware Algorithms, Systems, and Architectures (PASA): <http://www.cse.psu.edu/conferences/icpppasa2016>.

CONFERENCE VENUES

Lodging

Marriott Courtyard Philadelphia Downtown
21 North Juniper Street
Philadelphia, PA, 19107
Phone: 215-496-3200

Conference Site

Temple University City Center
1515 Market Street
Philadelphia, PA 19102

Reception

Estia Restaurant
1405-07 Locust Street
Philadelphia, PA 19102

Banquet

Joy Tsin Lau
1026-28 Race Street
Philadelphia, PA 19107