

Tentative workshop program

The Second IEEE International Workshop on Benchmarking, Performance Tuning and Optimization for Big Data Applications (BPOD), December 10
Workshop Chairs: Zhiyuan Chen, Jianwu Wang, Feng Chen, Yiming Ying

Time	Title	Presenter/Authors
8:30-8:35	Welcome, opening remark	Zhiyuan Chen & Feng Chen
Session 1: Workflow, Benchmark, Scientific Data		
8:35-8:55	Hoagie: A Database and Workload Generator using Published Specifications	Shahram Ghandeharizadeh, Haoyu Huang
8:55-9:20	Deep Learning for Enhancing Fault Tolerant Capabilities of Scientific Workflows	Alok Singh, Ilkay Altintas, Malachi Schram, Nathan Tallent
9:20-9:40	Blockchain Based Provenance Sharing of Scientific Workflows	Wanghu Chen, Xiaoyan Liang, Jing Li, Hongwu Qin, Yuxiang Mu, and Jianwu Wang
9:40-10:00	Toward Scalable Analysis of Multidimensional Scientific Data: A Case Study of Electrode Arrays	Ye Niu, Abdullah Al-Mamun, Hui Lin, Tonglin Li, Yi Zhao, and Dongfang Zhao
10:00-10:20	Coffee Break	
Session 2: Machine Learning & Analytics I		
10:20-10:45	A study of Exact Ridge Regression for Big Data	Wanchih Chiang, Xiang Liu, Tonglin Zhang, and Baijian Yang
10:45-11:10	A Flexible-blocking Based Approach for Performance Tuning of Matrix Multiplication Routines for Large Matrices with Edge Cases	Md Mosharaf Hossain, Thomas M. Hines, Sheikh Rabiul Islam, Sheikh K. Ghafoor, and Ramakrishnan Kannan
11:10-11:35	Performance and Memory Trade-offs of Deep Learning Object Detection in Fast Streaming High-Definition Images	Aishwarya Srivastava, Dung Nguyen, Siddhant Aggarwal, Andre Luckow, Edward Duffy, Ken Kennedy, Marcin Ziolkowski, and Amy Apon
11:35-12pm	City-Wide Influenza Forecasting based on Multi-Source Data	Kun Su, Yu Xiong, Li Qi, Yu Xia, Baisong Li, Lin Yang, Qin Li, Wenge Tang, Xian Li, Xiaowen Ruan, Shaofeng Lu, Xianxian Chen, Chaobo Shen, Boran Hong, Jiaying Xu, Liang Xu, Mei Han, and Jing Xiao
12-1:30	Lunch	
1:30-2:15	Invited Talk: Benchmarking for Big Data Applications with the DataBench Framework	Dr. Arne Berre
Session 3: Optimization and Tuning		
2:15-2:40	Enhancing the Scalability and Performance of Iterative Graph Algorithms on Apache Storm	Sachini Jayasekara, Shanika Karunasekera, and Aaron Harwood
2:40-3:05	One Self-Adaptive Memory Scheduling Algorithm for the Shuffle Process in Spark Platform	Jungang Xu, Shanshan Huang, Renfeng Liu, and Pengfei Li
3:05-3:30	Big data storage technologies: a case study for web-based LiDAR visualization	David Deibe, Margarita Amor, and Ramón Doall
3:30-3:50	Coffee Break	
Session 4: Machine Learning and Analytics II		
3:50-4:15	Predicting the Computational Cost of Deep Learning Models	Daniel Justus, John Brennan, Stephen Bonner, and Stephen McGough
4:15-4:40	Key based Deep Data Locality on Hadoop	Sungchul lee, Juyeon Jo, and Yoohwan Kim,
4:40-5:00	An Efficient Multi-objective Genetic Algorithm for Cloud Computing: NSGA-G	Trung-Dung Le, Verena Kantere, and Laurent d'Orazio
5:00-5:20	The Challenge of a Strong Speed-Up of a Bio-Medical Big Data Application	Marco Strutz, Bjoern Lindequist, Hermann Hessling, and Achim Streit
5:20-5:40	In-memory Blockchain: Toward Efficient and Trustworthy Data Provenance for HPC Systems	Abdullah Al-Mamun, Tonglin Li, Mohammad Sadoghi, and Dongfang Zhao
5:40-6:00	Achieving Horizontal Scalability in Density-based Clustering for URLs	Azadeh Faroughi, Reza Javidan, Marco Mellia, Andrea Morichetta, Francesca Soro, and Martino Trevisan
6-6:05	Closing remarks	Zhiyuan Chen & Feng Chen

Each long paper presentation is about 25 minutes (20 minutes talk + 5 minutes question answering).

Each short paper presentation is about 20 minutes (15 minutes talk + 5 minutes question answering).