

The Fourth IEEE International Workshop on Benchmarking, Performance Tuning and Optimization for Big Data Applications (BPOD), Thursday, December 10

Workshop Chairs: Zhiyuan Chen, Jianwu Wang, Feng Chen, Yiming Ying

Workshop website: <https://conferences.computer.org/ijmpas/BPOD/>

Time (US Eastern Time)	Title	Presenter/Author
Session 1 (Analytics & Optimization)		
9-9:05 am	Opening Remark	Workshop Chairs
9:05-9:25 am	S17212: Autoencoder-based outlier detection for sparse, high dimensional data	Wanghu Chen, Huijun Li, Jing Li, Ali Arshad
9:25-9:45 am	S17206: Impact of Map-Reduce framework on Hadoop and Spark MR Application Performance	Ishaan Lagwankar, Ananth Narayan Sankaranarayanan, Subramaniam Kalambur
9:45-10:00 am	S17207: A Topology-Aware Performance Prediction Model for Distributed Deep Learning on GPU Clusters	Zheyu Lin, Xukun Chen, Hanyu Zhao, Yunteng Luan, Zhi Yang, Yafei Dai,
10:00-10:15am	S17208: Benchmarking performance of RaySGD and Horovod for big data applications	Shruti Kunde, Amey Pandit, Rekha Singhal
Break		
Session 2 (Analytics & Optimization)		
10:30-10:50	S17209: Not Half Bad: Exploring Half-Precision in Graph Convolutional Neural Networks	John Brennan, Stephen Bonner, Amir Atapour-Abarghouei, Philip Jackson, Boguslaw Obara, Andrew Stephen McGough
10:50-11:10	S17203: Exploring Characteristics of Inter-cluster Machines and Cloud Applications on Google Clusters	Yuhui Lin, Adam Barker, Sheriffo Ceesay
11:10-11:30	BigD601: Making an Array Database Language Server-Side Extensible	Otoniel José Campos Escobar, Peter Baumann, Dimitar Misev
11:30-11:45	S17213: Towards Collaborative Optimization of Cluster Configurations for Distributed Dataflow Jobs	Jonathan Will, Jonathan Bader, Lauritz Thamsen
Lunch Break		
1-2 pm	S17216: Keynote: Benchmarks Big Data and HPC	Geoffrey Fox, Indiana University Bloomington
Session 3 (Benchmark)		
2:00-2:20 pm	BigD595: Evaluating NoSQL Systems for Decision Support: An Experimental Approach	Tomas Llano-Rios, Mohamed Khalefa, Antonio Badia
2:20-2:40 pm	S17204: Benchmarking HOAP for Scalable Document Data Management: A First Step	Yifan Tian, Michael Carey, Ian Maxon
2:40-2:55 pm	S17205: Performance Benchmarking of Automated Sentence Denoising using Deep Learning	Zhantong Liang, Abdou Youssef
2:55-3:10 pm	S17211: RDFINT: A Benchmark for Comparing Data Warehouse with Virtual Integration Approaches for Integration of RDF Data	Samson Oni, Kajal Pansare, Sukrit Arneja, Zhiyuan Chen, Adina Crainiceanu, Don Needham
3:10-3:25 pm	Break	
Session 4 (Analytics & Optimization)		
3:25-3:45 pm	BigD541: Computational performance of heterogeneous ensemble frameworks on high-performance computing platforms	Linhua Wang, Prem Timsina, Gaurav Pandey
3:45-4:05 pm	N246: Improving Join Reordering for Large Scale Distributed Computing	Amogh Margoor, Mayur Bhosale
4:05-4:20 pm	S17214: A Simple Low Cost Parallel Architecture for Big Data Analytics	Carlos Ordonez, Sikder Tahsin Al-Amin, and Xiantian Zhou
4:20-4:35 pm	S17210: Uncovering Performance Interference of Multi-Tenants in Big Data Environments	Adriano Lange, Tiago Rodrigo Kepe, Eduardo Cunha de Almeida, Marcos Sfair Sunye
4:35-4:50 pm	S17201: Automatic Tuning of Hyperparameters for Neural Networks in Serverless Cloud	Alex Kaplunovich, Yelena Yesha
Closing Remarks		