

The Third IEEE International Workshop on Benchmarking, Performance Tuning and Optimization for Big Data Applications (BPOD), Wednesday, December 11

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

Workshop website: <https://userpages.umbc.edu/~jianwu/BPOD>

Time	Title	Presenter/Authors
Session 1: Benchmark		
10:05-10:20	An Empirical Study of Rabin Fingerprinting Parameters	Paul Lu, Owen Randall, and Emma McDonald
10:20-10:40	Multidimensional Preference Query Optimization on Infrastructure Monitoring Systems	Yinghua Qin and Gheorghii Guzun
10:40-11:00	Performance Benchmarking of Data Argumentation and Deep Learning for Tornado Prediction	Carlos Barajas, Matthias Gobbert, and Jianwu Wang
11:00-11:20	Mix and Rank: A Framework for Benchmarking Recommender Systems	Bibek Paudel, Dragi Kocev, and Tome Eftimov
11:20-11:40	GeoYCSB: A Benchmark Framework for the Performance and Scalability Evaluation of NoSQL Databases for Geospatial Workloads	Suneuy Kim and Yuvraj Kanwar
11:40-11:55	Towards a High-Level Description for Generating Stream Processing Benchmark Applications	Alessio Pagliari, Fabrice Huet, and Guillaume Urvoy-Keller
11:55-12:10	Measuring, Quantifying, and Predicting the Cost-Accuracy Tradeoff	Matt Baughman, Nifesh Chakubaji, Kristis Kreics, Hong-Linh Truong, and Kyle Chard
12:10-2:30	Lunch and Main Conference Keynote	
Session 2: Analytics		
2:30-2:50	An Empirical Study of Quad-Level Cell (QLC) NAND Flash SSDs for Big Data Applications	Shuwen Liang, Zhi Qiao, Sihai Tang, Song Fu, and Weisong Shi
2:50-3:10	An Experimental Comparison of GPU Techniques for DBSCAN Clustering	Hamza Mustafa, Eleazar Leal, and Le Gruenwald
3:10-3:30	A Gray-box Testing Method for Divide & Conquer in Image Processing	Marco Strutz, Hermann Hessling, and Achim Streit
3:30-3:50	GHOSTZ PW/GF: distributed parallel homology search system for large-scale metagenomic analysis	Kenta Machida and Osamu Tatebe
3:50-4:10	Search for K: Assessing Five Topic Modeling Approaches to 120,000 Canadian Articles	Qiang Fu, Yufan Zhuang, Jiaxin Gu, Yushu Zhu, Huihui Qin, and Xin Guo
4:10-4:30	Coffee Break	
Session 3: Tuning		
4:30-4:50	Analysis and Prediction of Data Transfer Throughput for Data-Intensive Workloads	Devarshi Ghoshal, Kesheng Wu, Eric Pouyoul, and Erich Strohmaier
4:50-5:10	Fast Stochastic Block Partitioning using a Single Commodity Machine	Md Abdul Motaleb Faysal and Shaikh Arifuzzaman
5:10-5:30	Cluster-size optimization within a cloud-based ETL framework for Big Data	Eftim Zdravevski, Petre Lameski, Ace Dimitrievski, Marek Grzegorowski, and Cas Apanowicz
5:30-5:50	GraphOpt: a Framework for Automatic Parameters Tuning of Graph Processing Frameworks	Muaz Twaty and Amine Ghrab
5:50-6:10	Benchmarking the Discretisation Level of Continuous Attributes: Theoretical and Experimental Approaches	Wanghu Chen, Chao Wang, Jing Li, Bo Yang, Yang Liu, and Jianwu Wang
6:10-6:25	Reactive Microservices in Commodity Resources	Divya Goel and Amaresh Nayak