

Day 1 – December 2				
09:00-09:10	Opening			
09:10-09:55	Keynote 1: Resource Allocation and Consensus in Edge Blockchain Systems Yuanyuan Yang, IEEE Fellow, Stony Brook University, USA			
09:55-10:40	Keynote 2: On Optimal Partitioning and Scheduling of DNNs in Mobile Cloud Computing Jie Wu, AAAS Fellow, IEEE Fellow, Temple University, USA			
10:40-10:50	Virtual Break			
10:50-12:10	Session A1 (Algorithm Track)	Session A2: (Algorithm Track)	Session A3: (Algorithm Track)	
12:10-13:30	Virtual Lunch Break			
13:30-14:50	Session B1 (Architecture Track)	Session B2 (Architecture Track)	Session C1 (Application Track)	Session C2 (Application Track)
14:50-16:10	Session C3 (Application Track)	Session C4 (Application Track)	Session C5 (Application Track)	
16:20-18:00	Session D1 (Performance Track)	Session D2 (Performance Track)	Session E1 (Security Track)	Session E2 (Security Track)
Day 2 – December 3				
09:30-10:15	Keynote 3: Towards the Acceleration of Enterprise AI Hui Lei, IEEE Fellow, VP, Futurewei Technologies			
10:15-10:35	Award Session: Best Paper Award			
10:35-11:55	Session F1 (System Track)	Session F2 (System Track)	Session F3 (System Track)	
11:55-13:30	Virtual Lunch Break			
13:30-14:50	Session G1 (Invited Paper Track)	Session G2 (Invited Paper Track)	Session G3 (Invited Paper Track)	
14:50-15:00	Virtual Break			
15:00-17:20	Session H1 (Workshop BDS)	Session H2 (Workshop EISIA)	Session H3 (Workshop HMMECA)	
17:20-17:30	Closing			

***China Standard Time / Hong Kong Time (UTC+8)**

Program

December 2, 2020

09:00 AM - 09:10 AM (UTC+8) Opening Session

09:10 AM - 9:55 AM (UTC+8) Keynote 1

Keynote Title: Resource Allocation and Consensus in Edge Blockchain Systems

Yuanyuan Yang, IEEE Fellow, Stony Brook University, USA

09:55 AM - 10:40 AM (UTC+8) Keynote 2

Keynote Title: On Optimal Partitioning and Scheduling of DNNs in Mobile Cloud Computing

Jie Wu, AAAS Fellow, IEEE Fellow, Temple University, USA

10:40 AM - 10:50 AM (UTC+8) Break

10:50 AM - 11:50 AM (UTC+8)

Session A1: Learning Algorithm

LTG-LSM: The Optimal Structure in LSM-tree Combined with Reading Hotness

Jiaping Yu, Huahui Chen, Jiangbo Qian and Yihong Dong

Improved MapReduce Load Balancing through Distribution-Dependent Hash Function Optimization

Zafar Ahmad, Sharmila Duppala, Rezaul Chowdhury and Steven Skiena

Efficient Sparse-Dense Matrix-Matrix Multiplication on GPUs Using the Customized Sparse Storage Format

Shaohuai Shi, Qiang Wang and Xiaowen Chu

10:50 AM - 12:10 PM (UTC+8)

Session A2: Algorithms for Cloud Systems

Joint Service Placement and Request Scheduling for Multi-SP Mobile Edge Computing Network

Zhengwei Lei, Hongli Xu, Liusheng Huang and Zeyu Meng

A similarity clustering-based deduplication strategy in cloud storage systems

Saiqin Long, Zhetao Li, Zihao Liu, Qingyong Deng, Sangyoon Oh and Nobuyoshi Komuro

A Fair Task Assignment Strategy for Minimizing Cost in Mobile Crowdsensing

Yujun Liu, Yongjian Yang, En Wang, Wenbin Liu, Dongming Luan, Xiaoying Sun and Jie Wu

Communication-Aware Load Balancing of the LU Factorization over Heterogeneous Clusters

Lucas Leandro Nesi, Lucas Mello Schnorr and Arnaud Legrand

10:50 AM - 12:10 PM (UTC+8)

Session A3: Algorithms for Networks

An Efficient Work-Stealing Scheduler for Task Dependency Graph

Chun-Xun Lin, Tsung-Wei Huang and Martin D. F. Wong

LBNN: Perceiving the State Changes of a Core Telecommunications Network via Linear Bayesian Neural Network

Yanying Lin, Kejiang Ye, Ming Chen, Naitian Deng, Tailin Wu, and Cheng-Zhong Xu

A Method to Detecting Artifact Anomalies in A Microservice Architecture

Faisal Fahmi, Pei-Shu Huang and Feng-Jian Wang

Contention resolution on a restrained channel

Elijah Hradovich, Marek Klonowski and Dariusz R. Kowalski

01:30 PM - 02:30 PM (UTC+8)

Session B1: Performance Prediction and Optimization

Real-Time Scheduling and Analysis of OpenMP Programs with Spin Locks

He Du, Xu Jiang, Tao Yang, Mingsong Lv and Wang Yi

Predicting Performance Degradation on Adaptive Cache Replacement Policy

Yi Zhang, Ran Cui, Mingsong Lv, Chuanwen Li and Qingxu Deng

Making Inconsistent Components More Efficient For Hybrid B+Trees

Xiongxiong She, Chengliang Wang and Fenghua Tu

01:30 PM - 02:30 PM (UTC+8)

Session B2: Edge and Persistent Memory

XOR-Net: An Efficient Computation Pipeline for Binary Neural Network Inference on Edge Devices

Shien Zhu, Luan H. K. Duong, and Weichen Liu

Load Balance Aware Data Sharing Systems In Heterogeneous Edge Environment

Sheng Chen, Zheng Chen, Siyuan Gu, Baochao Chen, Junjie Xie and Deke Guo

Themis: Malicious Wear Detection and Defense for Persistent Memory File Systems

Wenbin Wang, Chaoshu Yang, Runyu Zhang, Shun Nie, Xianzhang Chen and Duo Liu

01:30 PM - 02:50 PM (UTC+8)

Session C1: Federated Learning and Reinforcement Learning

Incentive Mechanism Design for Federated Learning: A Two-stage Stackelberg Game Approach

Guiliang Xiao, Mingjun Xiao, Guoju Gao, Sheng Zhang, Hui Zhao and Xiang Zou

Time Efficient Federated Learning with Semi-asynchronous Communication

Jiangshan Hao, Yanchao Zhao and Jiale Zhang

Multi-agent Fault-tolerant Reinforcement Learning with Noisy Environments

Canhui Luo, Xuan Liu, Xinning Chen and Juan Luo

Decentralized Exploration of a Structured Environment Based on Multi-agent Deep Reinforcement Learning

Dingjie He, Dawei Feng, Hongda Jia and Hui Liu

01:30 PM - 02:50 PM (UTC+8)

Session C2: Crowd Sourcing and Parallel Acceleration

D2D-Enabled Reliable Data Collection for Mobile Crowd Sensing

Pengfei Wang, Zhen Yu, Chi Lin, Leyou Yang, Yaqing Hou and Qiang Zhang

Improving the Applicability of Visual Peer-to-Peer Navigation with Crowdsourcing

Erqun Dong, Jianzhe Liang, Zeyu Wang, Jingao Xu, Longfei Shangguan, Qiang Ma and Zheng Yang

Massively Parallel Causal Inference of Whole Brain Dynamics at Single Neuron Resolution

Wassapon Watanakesuntorn, Keichi Takahashi, Kohei Ichikawa, Joseph Park, George Sugihara, Ryousei Takano, Jason Haga and Gerald M. Pao

An Effective Design to Improve the Efficiency of DPUs on FPGA

Yutian Lei, Qingyong Deng, Saiqin Long, Shaohui Liu and Sangyoon Oh

02:50 PM - 04:10 PM (UTC+8)

Session C3: Localization and Cross-technology Communication

Lightweight Mobile Devices Indoor Location Based on Image Database

Ran Gao, Yanchao Zhao and Maoxing Tang

A Dynamic Escape Route Planning Method for Indoor Multi-floor Buildings Based on Real-time Fire Situation Awareness

Chun Wang, Juan Luo, Cuijun Zhang and Xuan Liu

Mitigating Cross-Technology Interference in Heterogeneous Wireless Networks based on Deep Learning

Weidong Zheng, Junmei Yao and Kaishun Wu

Accelerating PageRank in Shared-Memory for Efficient Social Network Graph Analytics

Baofu Huang, Zhidan Liu and Kaishun Wu

02:50 PM - 04:10 PM (UTC+8)

Session C4: Scheduling in Edge Environment

Using Configuration Semantic Features and Machine Learning Algorithms to Predict Build Result in Cloud-Based Container Environment

Yiwen Wu, Yang Zhang, Junsheng Chang, Bo Ding, Tao Wang and Huaimin Wang

Joint Service Placement and Computation Offloading in Mobile Edge Computing: An Auction-based Approach

Lei Zhang, Zhihao Qu, Baoliu Ye and Bin Tang

Multi-user Edge-assisted Video Analytics Task Offloading Game based on Deep Reinforcement Learning

Yu Chen, Sheng Zhang, Mingjun Xiao, Zhuzhong Qian, Jie Wu and Sanglu Lu

Accelerating Deep Learning Tasks with Optimized GPU-assisted Image Decoding

Lipeng Wang, Qiong Luo and Shengen Yan

02:50 PM - 04:10 PM (UTC+8)

Session C5: Application and Security

Scheduling Rechargeable UAVs for Long Time Barrier Coverage

Zhouqing Han, Xiaojun Zhu and Lijie Xu

Use of Genetic Programming Operators in Data Replication and Fault Tolerance

Syed Mohtashim Abbas Bokhari and Oliver Theel

Multiple Balance Subsets Stacking for Imbalanced Healthcare Datasets

Yachao Shao, Tao Zhao, Xiaoning Wang, Xiaofeng Zou and Xiaoming Fu

Securing App Behaviors in Smart Home: A Human-App Interaction Perspective

Jinlei Li, Yan Meng, Lu Zhou and Haojin Zhu

04:20 PM - 05:40 PM (UTC+8)

Session D1: Edge-Cloud Performance

Task Offloading in Trusted Execution Environment empowered Edge Computing

Yuepeng Li, Deze Zeng, Lin Gu, Andong Zhu and Quan Chen

Gecko: Guaranteeing Latency SLO on a Multi-Tenant Distributed Storage System

Zhenyu Leng, Dejun Jiang, Liuying Ma and Jin Xiong

A Learning-based Dynamic Load Balancing Approach for Microservice Systems in Multi-cloud Environment

Jieqi Cui, Pengfei Chen and Guangba Yu

<p>Enhancing Availability for the MEC Service: CVaR-based Computation Offloading Shengli Pan, Zhiyong Zhang, Tao Xue and Guangmin Hu</p>
<p>04:20 PM - 05:20 PM (UTC+8) Session D2: Performance of Distributed Systems</p>
<p>Intermediate Value Size Aware Coded MapReduce Yamei Dong, Bin Tang, Baoliu Ye, Zhihao Qu and Sanglu Lu</p> <p>A Customized Reinforcement Learning based Binary Offloading in Edge Cloud Yuepeng Li, Lvhao Chen, Deze Zeng and Lin Gu</p> <p>Optimal Use Of The TCP/IP Stack in User-space Storage Applications with ADQ feature in NIC Ziye Yang, Ben Walker, James R Harris, Yadong Li, and Gang Cao</p>
<p>04:20 PM - 05:40 PM (UTC+8) Session E1: Cyber Physical Security</p>
<p>SIC²: Securing Microcontroller Based IoT Devices with Low-cost Crypto Coprocessors Bryan Pearson, Cliff Zou, Yue Zhang, Zhen Ling and Xinwen Fu</p> <p>Intelligent detection algorithm against UAVs' GPS spoofing attack Shenqing Wang, Jiang Wang, Chunhua Su, and Xinshu Ma</p> <p>An Efficient and Scalable Sparse Polynomial Multiplication Accelerator for LAC on FPGA Jipeng Zhang, Zhe Liu, Hao Yang, Junhao Huang and Weibin Wu</p> <p>Secure and Verifiable Data Access Control Scheme With Policy Update and Computation Outsourcing for Edge Computing Yue Guan, Songtao Guo, Pan Li and Yuanyuan Yang</p>
<p>04:20 PM - 06:00 PM (UTC+8) Session E2: AI and Distributed System Security</p>
<p>Secure Door on Cloud: A Secure Data Transmission Scheme to Protect Kafka's Data Hanyi Zhang, Liming Fang, Keyu Jiang, Weiting Zhang, Minghui Li and Lu Zhou</p> <p>A Solution to Data Accessibility Across Heterogeneous Blockchains Zhihui Wu, Yang Xiao, Enyuan Zhou, Qingqi Pei, and Quan Wang</p> <p>PrivAG: Analyzing Attributed Graph Data with Local Differential Privacy Zichun Liu, Liusheng Huang, Hongli Xu, Wei Yang and Shaowei Wang</p> <p>Exploring Data Correlation between Feature Pairs for Generating Constraint-based Adversarial Examples Yunzhe Tian, Yingdi Wang, Endong Tong, Wenjia Niu, Liang Chang, Qi Alfred Chen, Gang Li and Jiqiang Liu</p>

A Deep Learning Framework Supporting Model Ownership Protection and Traitor Tracing
Guowen Xu, Hongwei Li, Yuan Zhang, Xiaodong Lin, Robert H. Deng and Xuemin Shen

December 3, 2020

09:30 AM - 10:15 AM (UTC+8) Keynote 3

Keynote Title: Towards the Acceleration of Enterprise AI

Hui Lei, IEEE Fellow, VP, Futurewei Technologies

10:15 AM - 10:35 AM (UTC+8) Award Session: Best Paper Award

10:35 AM - 11:55 AM (UTC+8)

Session F1: Data Processing

A social link based private storage cloud

Michalis Konstantopoulos, Nikos Chondros and Mema Roussopoulos

Enabling Generic Verifiable Aggregate Query on Blockchain Systems

Yanchao Zhu, Zhao Zhang, Cheqing Jin, and Aoying Zhou

SrSpark: Skew-resilient Spark based on Adaptive Parallel Processing

Yijie Shen, Jin Xiong and Dejun Jiang

Optimizing Multi-way Theta Join for Data Skew in Sub-Second Stream Computing

Xiaopeng Fan, Xinchun Liu, Yang Wang, Youjun Wang, and Jing Li

10:35 AM - 11:55 AM (UTC+8)

Session F2: Resource and Data Management

OOOPS: An Innovative Tool for IO Workload Management on Supercomputers

Lei Huang and Si Liu

URFS: A User-space Raw File System based on NVMe SSD

Yaofeng Tu, Yinjun Han, Zhenghua Chen, Zhengguang Chen and Bing Chen

DyRAC: Cost-aware Resource Assignment and Provider Selection for Dynamic Cloud Workloads

Yannis Sfakianakis, Manolis Marazakis and Angelos Bilas

WMAlloc: A Wear-Leveling-Aware Multi-Grained Allocator for Persistent Memory File Systems

Shun Nie, Chaoshu Yang, Runyu Zhang, Wenbin Wang, Duo Liu and Xianzhang Chen

10:35 AM - 11:55 AM (UTC+8)

Session F3: Secure and Reliable Systems

Optimizing Complex OpenCL Code for FPGA: A Case Study on Finite Automata Traversal
Marziyeh Nourian, Mostafa Eghbali Zarch and Michela Becchi

FastCredit: Expediting Credit-based Proactive Transports in Datacenters
Dezun Dong, Shan Huang, Zejia Zhou, Wenxiang Yang and Hanyi Shi

Proactive Failure Recovery for Stateful NFV
Zhenyi Huang and Huawei Huang

TEEp: Supporting Secure Parallel Processing in ARM TrustZone
Zinan Li, Wenhao Li, Yubin Xia and Binyu Zang

01:30 PM - 02:50 PM (UTC+8)

Session G1: Scheduling and Resource Management

Compressive Sensing based Predictive Online Scheduling with Task Colocation in Cloud Data Center
Yunhin Chan, Ke Luo and Xu Chen

Non-Technical Losses Detection in Smart Grids: An Ensemble Data-Driven Approach
Yufeng Xing, Lei Guo, Zongchao Xie, Lei Cui, Longxiang Gao and Shui Yu

Virtual Machine Consolidation for NUMA Systems: A Hybrid Heuristic Grey Wolf Approach
Kangli Hu, Weiwei Lin, Tiansheng Huang, Keqin Li and Like Ma

A Novel Classification Model to Predict Batch Job Failures in Co-located Cloud
Yurui Li, Weiwei Lin, Keqin Li, James Z. Wang, Fagui Liu and Jie Liu

01:30 PM - 02:50 PM (UTC+8)

Session G2: Distributed System Design and Implementation

ABC: An Auction-Based Blockchain Consensus-Incentive Mechanism
Zhengpeng Ai, Yuan Liu and Xingwei Wang

A Trustworthy Blockchain-based Decentralised Resource Management System in the Cloud
Zhiming Zhao, Chunming Rong and Martin Gilje Jaatun

DCVP: Distributed Collaborative Video Stream Processing in Edge Computing
Shijing Yuan, Jie Li, Chentao Wu, Yusheng Ji and Yongbing Zhang

Efficient Post-quantum Identity-based Encryption with Equality Test
Willy Susilo, Dung Hoang Duong and Huy Quoc Le

01:30 PM - 02:30 PM (UTC+8)

Session G3: Federated Learning and Deep Learning

Robust Federated Learning Approach for Travel Mode Identification from Non-IID GPS Trajectories

Yuanshao Zhu, Shuyu Zhang, Yi Liu, Dusit Niyato and James J.Q. Yu

Deep Spatio-Temporal Attention Model for Grain Storage Temperature Forecasting

Shanshan Duan, Weidong Yang, Xuyu Wang, Shiwen Mao and Yuan Zhang

Proactive Content Caching for Internet-of-Vehicles based on Peer-to-Peer Federated Learning

Zhengxin Yu, Jia Hu, Geyong Min, Han Xu and Jed Mills

03:00 PM - 04:20 PM (UTC+8)

Session H1: Workshop - Big Data Systems

How Fast Can We Insert? An Empirical Performance Evaluation of Apache Kafka

Guenter Hesse, Christoph Matthies and Matthias Uflacker

Performance Modeling and Tuning for DFT Calculations on Heterogeneous Architectures

Hadia Ahmed, David Williams-Young, Khaled Z. Ibrahim and Chao Yang

Graph-based Approaches for the Interval Scheduling Problem

Panagiotis Oikonomou, Nikos Tziritas, Georgios Theodoropoulos, Maria Koziri, Thanasis Loukopoulos and Samee U. Khan

MPI parallelization of NEUROiD models using docker swarm

Raghu Sesha Iyengar and Mohan Raghavan

03:00 PM - 05:20 PM (UTC+8)

Session H2: Workshop - Edge Intelligence for Smart IoT Applications

S-GAT: Accelerating Graph Attention Networks Inference on FPGA Platform with Shift Operation

Weian Yan, Weiqin Tong, and Xiaoli Zhi

Explainable Congestion Attack Prediction and Software-level Reinforcement in Intelligent Traffic Signal System

Xiaojin Wang, Yingxiao Xiang, Wenjia Niu, Endong Tong, and Jiqiang Liu

Dynamic-Static-based Spatiotemporal Multi-Graph Neural Networks for Passenger Flow Prediction

Jingyan Ma, Jingjing Gu, Qiang Zhou, Qihong Wang and Ming Sun

Incentive-driven Data Offloading and Caching Replacement Scheme in Opportunistic Mobile Networks

Tong Wu, Xuxun Liu, Deze Zeng, Huan Zhou and Shouzhi Xu

Adaptive DNN Partition in Edge Computing Environments

Weiwei Miao, Zeng Zeng, Lei Wei, Shihao Li, Chengling Jiang and Zhen Zhang

Efficient Edge Service Migration in Mobile Edge Computing

Zeng Zeng, Shihao Li, Weiwei Miao, Lei Wei, Chengling Jiang, Chuanjun Wang and Mingxuan Zhang

A protocol-independent container network observability analysis system based on eBPF

Chang Liu, Zhengong Cai, Bingshen Wang, Zhimin Tang, and Jiaxu Liu

03:00 PM - 04:40 PM (UTC+8)

Session H3: Workshop - Heterogeneous Multi-access Mobile Edge Computing and Applications

Performance Guaranteed Single Link Failure Recovery in SDN Overlay Networks

Lilei Zheng, Hongli Xu, Suo Chen and Liusheng Huang

A Personal Distributed Real-time Collaborative System

Michalis Konstantopoulos, Nikos Chondros and Mema Roussopoulos

Cooperative Resource Sharing Strategy With eMBB Cellular and C-V2X Slices

Yan Liang, Xin Chen, Shuang Chen and Ying Chen

CP-BZD Repair Codes Design for Distributed Edge Computing

Shuangshuang Lu, Chanting Zhang and Mingjun Dai

Computation Task Scheduling and Offloading Optimization for Collaborative Mobile Edge Computing

Bin Lin, Xiaohui Lin, Shengli Zhang, Hui Wang and Suzhi Bi

05: 20 PM - 05:30 PM (UTC+8) Closing Session