

NAS 2017 PROGRAM At-A-GLANCE

August 6 (Sunday)	
20:00 – 21:00	Registration open (China merchant)

August 7 (Monday)	
07:00	Registration open (China merchant) Continental Breakfast (China merchant)
08:30 - 09:00	Opening Remarks (China merchant)
09:00 - 10:00	Keynote 1 (China merchant) Title: Internet of Vehicles: From Intelligent Grid to Autonomous Cars Speaker: Prof. Mario Gerla, University of California, Los Angeles
10:00 - 10:20	Coffee break
10:20 – 12:00	Session 1: Architecture (China merchant)
12:00 - 13:30	Lunch break
13:30 - 14:50	Session 2: Networking (China merchant)
14:50 - 16:30	Session 3: Storage 1 (China merchant)
16:30 - 16:50	Coffee break
16:50 - 18:10	Session 4: Short paper 1 (China merchant)
18:30 - 20:30	Banquet

August 8 (Tuesday)	
08:00	Registration open (China merchant) Continental Breakfast (China merchant)
09:00 - 10:00	Keynote 2 (China merchant) Title: Paving the Way towards NFV: An Infrastructure based Approach Speaker: Prof. Tao Li, University of Florida, NSF Program Director
10:00 - 10:20	Coffee break
10:20 – 12:00	Session 5:Storage 2 (China merchant)
12:00 - 13:30	Lunch break
13:30 - 14:50	Session 6: Architecture and Networking (China merchant)
14:50 - 16:10	Session 7: Storage 3 (China merchant)
16:10 - 16:30	Coffee break
16:30 - 18:00	Session 8: Short paper 2 (China merchant)
18:00 - 20:00	Poster session

August 7 (Monday)

Session 1 Architecture (10:20 – 12:00)

Chair: Hamid Sarbazi-Azad, Sharif University of Technology and IPM azad@sharif.edu

DD-L1D: Improving the Decoupled L1D Efficiency for GPU architecture
Weiguang Yang, Yuxin Wang, Yulong Yu, Guangyuan Kan and He Guo

Data Block Partitioning for Recovering Stuck-at Faults in PCMs
Marjan Asadinia, Majid Jalili and Hamid Sarbazi-Azad

ALARM: A Location-Aware Redistribution Method to Improve 3D FG NAND Flash Reliability
Yue Zhu, Fei Wu, Qin Xiong, Zhonghai Lu and Changsheng Xie

Multicast Scheduling with Markov Chains in Fat-tree Data Center Networks
Guozhi Li, Songtao Guo, Guiyan Liu and Yuanyuan Yang

Session 2 Network (13:30 - 14:50)

Chair: Prof. Ghasem Mirjalily, Shenzhen Research Institute of Big Data g.mirjalily@gmail.com

Enhancing Next Generation Passive Optical Network Stage 2 (NG-PON2) with Channel Bonding
Liang Zhang, Yuanqiu Luo, Nirwan Ansari, Bo Gao, Xiang Liu and Frank Effenberger

Contact Duration-Aware Routing in Delay Tolerant Networks
Tuan Le and Mario Gerla

MTM: A Reliable Multiple Trees Multicast for Data Center Network
Xin Xiong and Tan Chen

Session 3 Storage 1 (14:50 - 16:30)

Chair: Ming Zhao, Arizona State University mingzhao@asu.edu

Megalloc*: Fast Distributed Memory Allocator for NVM-based Cluster
Songping Yu, Nong Xiao and Mingzhu Deng

TraceRAR: An I/O Performance Evaluation Tool for Replaying, Analyzing, and Regenerating Traces
Bingzhe Li, Farnaz Toussi, Clark Anderson, David Lilja and David Du

A High-Performance Persistent Identifier Management Protocol
Fatih Berber and Ramin Yahyapour

Performance Evaluation and Modeling of HPC I/O on Non-Volatile Memory
Wei Liu, Kai Wu, Jialin Liu, Feng Chen and Dong Li

Session 4 Short paper 1 (16:50 - 18:10)

Chair: Kejiang Ye, Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences

kj.ye@siat.ac.cn

A-MapCG: An Adaptive MapReduce Framework for GPUs
Lifeng Liu, Yue Zhang, Meilin Liu, Wang Chongjun and Jun Wang

Performance Tuning for Big Data Applications in Docker Containers
Kejiang Ye and Yunjie Ji

An Experimental Study on Deep Learning Based on Different Hardware Configurations
Chen Zhang and Qiang Cao

Automatic Collector for Dynamic Cloud Performance Information
Olaf Elzinga, Spiros Koulouzis, Arie Taal and Zhiming Zhao

Facilitating Workload Aware Storage Platform by using Machine Learning Technics
Wubin Li, Fereydoun Farrahi Moghaddam, Parisa Heidari, Yves Lemieux and Abdelouahed Gherbi

SpyStorage: A Highly Reliable Multi-Cloud Storage with Secure and Anonymous Data Sharing
Pengyan Shen, Wan Liu, Zheng Wu, Mingzhong Xiao and Quanqing Xu

August 8 (Tuesday)

Session 5 Storage 2 (10:20 – 12:00)

Chair: Yuanqiu Luo, Huawei Technologies USA yuanqiu.luo@huawei.com

Service Migrations in the Cloud for Mobile Accesses: A Reinforcement Learning Approach
Shan Cao, Chengzhong Xu and Yang Wang

WPS: A Workload-aware Placement Scheme for Erasure-Coded In-Memory Stores
Shuang Wang, Jianzhong Huang, Xiao Qin and Changsheng Xie

Reducing Chunk Fragmentation for In-line Delta Compressed and Deduplicated Backup Systems
Yucheng Zhang, Dan Feng, Yu Hua, Yuchong Hu, Wen Xia, Min Fu, Xiaolan Tang, Zhikun Wang, Fangting Huang and Yukun Zhou

Analysis and Correlation of Application I/O Performance and System-Wide I/O Activity
Sandeep Madireddy, Prasanna Balaprakash, Philip Carns, Robert Latham, Robert Ross, Shane Snyder and Stefan Wild

Session 6 Architecture and Networking (13:30 - 14:50)

Chair: Yifeng Zhu, University of Maine Yifeng.Zhu@maine.edu

Optimizing Energy Consumption on HPC Systems with a Multi-level Checkpointing Mechanism
Muhammad Alfian Amrizal and Hiroyuki Takizawa

Towards Robust and Accurate Similar Trajectory Discovery: Weak-parametric Approaches
Yupeng Tuo

Detecting Information Theft Based on Mobile Network Flows for Android Users
Zhenyu Cheng, Xunxun Chen, Yongzheng Zhang, Shuhao Li and Yafei Sang

Session 7 Storage 3 (14:50 - 16:10)

Chair: Song Jiang, University of Texas, Arlington song.jiang@uta.edu

Extending Lifetime of SSD in RAID5 Systems through a Reliable Hierarchical Cache
RuiYe, Wentao Meng and Shenggang Wan

ISM - An Intra-Stripe Data Migration Approach for RAID-5 Scaling
Jie Liang, Yinlong Xu, Yongkun Li and Yubiao Pan

Kaleido: Enabling Efficient Scientific Data Processing on Big-Data Systems
Saman Biokaghazadeh, Shujia Zhou and Ming Zhao

Session 8 Short paper 2 (16:30 - 18:00)

Chair: Yinjin Fu, PLA University of Science and Technology yinjinfu@gmail.com

A Write-Through Cache Method to Improve Small Write Performance of SSD-based RAID
Linjun Mei, Dan Feng, Lingfang Zeng, Jianxi Chen and Jingning Liu

DualStack: A High Efficient Dynamic Page Scheduling Scheme in Hybrid Main Memory
Zhen Zhang, Yinjin Fu and Guyu Hu

A Hash-based Space-efficient Page-level FTL for Large-capacity SSDs
Fan Ni, Chunyi Liu, Yang Wang, Chengzhong Xu, Xiao Zhang and Song Jiang

Revisiting Updating Schemes for Erasure-coded In-Memory Stores
Jie Xia, Jianzhong Huang, Xiao Qin, Qiang Cao and Changsheng Xie

Exploiting Virtual Metadata Servers to Provide Multi-level Consistency for Key-value Object-based Data Store
Xiaozhao Zhuang, Xiaoyang Qu, Zhiyong Lu, Jiguang Wan and Changsheng Xie

Binary Index and Journal Embedding in The Linear Tape File System
Klaus Birkelund Jensen and Brian Vinter

Balancing the Storage in a Deduplication Cluster
Giacomo Grangia, Quanqing Xu, Andrea Bianco and Paolo Giaccone

Poster Session (18:00 – 20:00)

SimpleBP: A Lightweight Branch Prediction Simulator for Effective Design Exploration
Chaobing Zhou, Libo Huang, Zhisheng Li, and Qiang Dou

hpFog: A FPGA-based Fog Computing Platform
Tze Hon Tan, Chia Yee Ooi, and M. N. Marsono

Development of Network Simulator for LWA/LAA Implementations
Kun-Hee Lee, Min-Taek Choi, Sung-Chul Choi, and Jae-Hoon Kim

Design of Universal Broker Architecture for Edge Networking
Joo-Hyun Kim, Seo-Hee Hong, So-Hyun Yang, and Jae-Hoon Kim

Performance Optimization of In-Memory File System in Distributed Storage System
Zhaowei Li, Yunlong Yan, Jintao Mo, Zhaocong Wen, and Junmin Wu

Branch Prediction Migration for Multi-core Architectures
Tan Zhang, Chaobing Zhou, Libo Huang, and Nong Xiao

FlashStorageSim: Performance Modeling for SSD Architectures
Krishna T. Malladi, Mu-Tien Chang, Dimin Niu, and Hongzhong Zheng

Rack Level Scheduling for Containerized Workloads
Qiumin Xu, Krishna T. Malladi, and Manu Awasthi

Privacy-Protected Data Collection in Wireless Medical Sensor Networks
Md Zakirul Alam Bhuiyan, Mdaliuz Zaman, Guojun Wang, Tian Wang, and Jie Wu

HealthEdge: Task Scheduling for Edge Computing with Health Emergency and Human Behavior
Consideration in Smart Homes
Haoyu Wang, Jiaqi Gong, Yan Zhuang, Haiying Shen, and John Lach