

# **2019 35th Symposium on Mass Storage Systems and Technologies (MSST 2019)**

**Santa Clara, California, USA  
20 – 24 May 2019**



**IEEE Catalog Number: CFP19257-POD  
ISBN: 978-1-7281-3921-0**

**Copyright © 2019 by the Institute of Electrical and Electronics Engineers, Inc.  
All Rights Reserved**

*Copyright and Reprint Permissions:* Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limit of U.S. copyright law for private use of patrons those articles in this volume that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

For other copying, reprint or republication permission, write to IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08854. All rights reserved.

***\*\*\* This is a print representation of what appears in the IEEE Digital Library. Some format issues inherent in the e-media version may also appear in this print version.***

IEEE Catalog Number:	CFP19257-POD
ISBN (Print-On-Demand):	978-1-7281-3921-0
ISBN (Online):	978-1-7281-3920-3
ISSN:	2160-195X

**Additional Copies of This Publication Are Available From:**

Curran Associates, Inc  
57 Morehouse Lane  
Red Hook, NY 12571 USA  
Phone: (845) 758-0400  
Fax: (845) 758-2633  
E-mail: [curran@proceedings.com](mailto:curran@proceedings.com)  
Web: [www.proceedings.com](http://www.proceedings.com)

CURRAN ASSOCIATES INC.  
**proceedings**  
.com

# 2019 35th Symposium on Mass Storage Systems and Technologies (MSST) MSST 2019

## Table of Contents

Program Committee ix.....

### Performance

Fighting with Unknowns: Estimating the Performance of Scalable Distributed Storage Systems with Minimal Measurement Data .1.....  
*Moo-Ryong Ra (AT&T Labs Research) and Hee Won Lee (AT&T Labs Research)*

A Performance Study of Lustre File System Checker: Bottlenecks and Potentials .7.....  
*Dong Dai (University of North Carolina at Charlotte), Om Rameshwar  
Gatla (Iowa State University), and Mai Zheng (Iowa State University)*

Scalable QoS for Distributed Storage Clusters using Dynamic Token Allocation .14.....  
*Yuhan Peng (Rice University), Qingyue Liu (Rice University), and Peter  
Varman (Rice University)*

FastBuild: Accelerating Docker Image Building for Efficient Development and Deployment of Container .28....  
*Zhuo Huang (Huazhong University of Science and Technology), Song Wu  
(Huazhong University of Science and Technology), Song Jiang (The  
University of Texas at Arlington), and Hai Jin (Huazhong University of  
Science and Technology)*

BFO: Batch-File Operations on Massive Files for Consistent Performance Improvement .38.....  
*Yang Yang (Huazhong University of Science and Technology), Qiang Cao  
(Huazhong University of Science and Technology), Hong Jiang (The  
University of Texas at Arlington), Li Yang (Huazhong University of  
Science and Technology), Jie Yao (Huazhong University of Science and  
Technology), Yuanyuan Dong (Alibaba Group), and Puyuan Yang (Alibaba  
Group)*

### High Performance Computing

vPFS+: Managing I/O Performance for Diverse HPC Applications .51.....  
*Ming Zhao (Arizona State University) and Yiqi Xu (VMware)*

Accelerating Relative-error Bounded Lossy Compression for HPC datasets with Precomputation-Based Mechanisms .65.....	
	<i>Xiangyu Zou (Harbin Institute of Technology, Shenzhen), Tao Lu (Marvell Technology Group), Wen Xia (Harbin Institute of Technology, Shenzhen), Xuan Wang (Harbin Institute of Technology, Shenzhen), Weizhe Zhang (Harbin Institute of Technology, Shenzhen), Sheng Di (Argonne National Laboratory), Dingwen Tao (University of Alabama), and Franck Cappello (Argonne National Laboratory)</i>
Efficient Encoding and Reconstruction of HPC Datasets for Checkpoint/Restart .79.....	
	<i>Jialing Zhang (University of Massachusetts Lowell), Xiaoyan Zhuo (University of Massachusetts Lowell), Aekyeung Moon (University of Massachusetts Lowell), Hang Liu (University of Massachusetts Lowell), and Seung Woo Son (University of Massachusetts Lowell)</i>

## Persistent RAM

Tiered-ReRAM: A Low Latency and Energy Efficient TLC Crossbar ReRAM Architecture .92.....	
	<i>Yang Zhang (Huazhong University of Science and Technology), Dan Feng (Huazhong University of Science and Technology), Wei Tong (Huazhong University of Science and Technology), Jingning Liu (Huazhong University of Science and Technology), Chengning Wang (Huazhong University of Science and Technology), and Jie Xu (Huazhong University of Science and Technology)</i>
vNVML: An Efficient User Space Library for Virtualizing and Sharing Non-Volatile Memories .103.....	
	<i>Chih Chieh Chou (Texas A&amp;M University), Jaemin Jung (Samsung Semiconductor), A. L. Narasimha Reddy (Texas A&amp;M University), Paul V. Gratz (Texas A&amp;M University), and Doug Voigt (Hewlett Packard Enterprise)</i>
Towards Virtual Machine Image Management for Persistent Memory .116.....	
	<i>Jiachen Zhang (Nankai University), Lixiao Cui (Nankai University), Peng Li (Nankai University), Xiaoguang Liu (Nankai University), and Gang Wang (Nankai University)</i>
Pattern-based Write Scheduling and Read Balance-oriented Wear-Leveling for Solid State Drivers .126.....	
	<i>Jun Li (Southwest University), Xiaofei Xu (Southwest University of China), Xiaoning Peng (Huaihua University), and Jianwei Liao (Southwest University of China)</i>
When NVMe over Fabrics Meets Arm: Performance and Implications .134.....	
	<i>Yichen Jia (Louisiana State University), Eric Anger (Arm Inc), and Feng Chen (Louisiana State University)</i>
Long-Term JPEG Data Protection and Recovery for NAND Flash-Based Solid-State Storage .141.....	
	<i>Yu-Chun Kuo (National Tsing Hua University), Ruei-Fong Chiu (National Tsing Hua University), and Ren-Shuo Liu (National Tsing Hua University)</i>
Wear-aware Memory Management Scheme for Balancing Lifetime and Performance of Multiple NVM Slots .148	
	<i>Chunhua Xiao (Chongqing University), Linfeng Cheng (Chongqing University), Lei Zhang (Chongqing University), Duo Liu (Chongqing University), and Weichen Liu (Nanyang Technological University)</i>

CeSR: A Cell State Remapping Strategy to Reduce Raw Bit Error Rate of MLC NAND Flash .161.....  
*Yutong Zhao (Wuhan National Lab for Optoelectronics), Wei Tong (Wuhan National Lab for Optoelectronics), Jingning Liu (Wuhan National Lab for Optoelectronics), Dan Feng (Wuhan National Lab for Optoelectronics), and Hongwei Qin (Wuhan National Lab for Optoelectronics)*

Parallel all the time: Plane Level Parallelism Exploration for High Performance SSDs .172.....  
*Congming Gao (East China Normal University; Chongqing University; University of Pittsburgh), Liang Shi (East China Normal University), Chun Jason Xue (City University of Hong Kong), Cheng Ji (City University of Hong Kong), Jun Yang (University of Pittsburgh), and Youtao Zhang (University of Pittsburgh)*

## Disk

Economics of Information Storage: The Value in Storing the Long Tail .185.....  
*James P Hughes (UCSC)*

DFPE: Explaining Predictive Models for Disk Failure Prediction .193.....  
*Yanwen Xie (Huazhong University of Science and Technology), Dan Feng (Huazhong University of Science and Technology), Fang Wang (Huazhong University of Science and Technology), Xuehai Tang (Chinese Academy of Sciences), Jizhong Han (Chinese Academy of Sciences), and Xinyan Zhang (Huazhong University of Science and Technology)*

Mitigate HDD Fail-Slow by Pro-actively Utilizing System-level Data Redundancy with Enhanced HDD Controllability and Observability .205.....  
*Jingpeng Hao (Rensselaer Polytechnic Institute), Yin Li (Rensselaer Polytechnic Institute), Xubin Chen (Rensselaer Polytechnic Institute), and Tong Zhang (Rensselaer Polytechnic Institute)*

## Coding

Adjustable Flat Layouts for Two-Failure Tolerant Storage Systems .217.....  
*Thomas Schwarz (Marquette University)*

AZ-Code: An Efficient Availability Zone Level Erasure Code to Provide High Fault Tolerance in Cloud Storage Systems .230.....  
*Xin Xie (Shanghai Jiao Tong University), Chentao Wu (Shanghai Jiao Tong University), Junqing Gu (Shanghai Jiao Tong University), Han Qiu (Shanghai Jiao Tong University), Jie Li (Shanghai Jiao Tong University), Minyi Guo (Shanghai Jiao Tong University), Xubin He (Temple University), Yuanyuan Dong (Alibaba Group), and Yafei Zhao (Alibaba Group)*

XORInc: Optimizing Data Repair and Update for Erasure-Coded Systems with XOR-Based In-Network Computation .244.....  
*Fang Wang (Huazhong University of Science and Technology), Yingjie Tang (Huazhong University of Science and Technology), Yanwen Xie (Huazhong University of Science and Technology), and Xuehai Tang (Chinese Academy of Sciences)*

Parity-Only Caching for Robust Straggler Tolerance .257.....  
*Mi Zhang (The Chinese University of Hong Kong), Qiuping Wang (The Chinese University of Hong Kong), Zhirong Shen (The Chinese University of Hong Kong), and Patrick P. C. Lee (The Chinese University of Hong Kong)*

**Deduplication**

Metademap: Deduplicating Metadata in Encrypted Deduplication via Indirection .269.....  
*Jingwei Li (University of Electronic Science and Technology of China), Patrick P. C. Lee (The Chinese University of Hong Kong), Yanjing Ren (University of Electronic Science and Technology of China), and Xiaosong Zhang (University of Electronic Science and Technology of China)*

CDAC: Content-Driven Deduplication-Aware Storage Cache .282.....  
*Yujuan Tan (Chongqing University), Jing Xie (Chongqing Univeristy), Congcong Xu (Chongqing University), Zhichao Yan (HPE), Hong Jiang (University of Texas Arlington), Yajun Zhao (Sangfor), Min Fu (Sangfor), Xianzhang Chen (Chongqing University), Duo Liu (Chongqing University), and Wen Xia (Harbin Institute of Technology)*

SES-Dedup: a Case for Low-Cost ECC-based SSD Deduplication .292.....  
*Zhichao Yan (University of Texas at Arlington), Hong Jiang (University of Texas at Arlington), Song Jiang (University of Texas at Arlington), Yujuan Tan (Chongqing University), and Hao Luo (Twitter)*

LIPA: A Learning-based Indexing and Prefetching Approach for Data Deduplication .299.....  
*Guangping Xu (Tianjin University of Technology), Bo Tang (Tianjin University of Technology), Hongli Lu (Tianjin University of Technology), Quan Yu (Wuhan University of Technology), and Chi Wan Sung (City University of Hong Kong)*

**Author Index 311** .....