

FAST '19: 17th USENIX Conference on File and Storage Technologies
February 25–28, 2019
Boston, MA, USA

Persistent Memory Systems

Reaping the performance of fast NVM storage with uDepot1
Kornilios Kourtis, Nikolas Ioannou, and Ioannis Koltsidas, *IBM Research*

Optimizing Systems for Byte-Addressable NVM by Reducing Bit Flipping17
Daniel Bittman, Darrell D. E. Long, Peter Alvaro, and Ethan L. Miller, *UC Santa Cruz*

Write-Optimized Dynamic Hashing for Persistent Memory31
Moohyeon Nam, *UNIST (Ulsan National Institute of Science and Technology)*; Hokeun Cha, *Sungkyunkwan University*;
Young-ri Choi and Sam H. Noh, *UNIST (Ulsan National Institute of Science and Technology)*; Beomseok Nam,
Sungkyunkwan University

Software Wear Management for Persistent Memories45
Vaibhav Gogte, *University of Michigan*; William Wang and Stephan Diestelhorst, *ARM*; Aasheesh Kolli, *Pennsylvania State University and VMware Research*; Peter M. Chen, Satish Narayanasamy, and Thomas F. Wenisch, *University of Michigan*

File Systems

Storage Gardening: Using a Virtualization Layer for Efficient Defragmentation in the WAFL File System65
Ram Kesavan, Matthew Curtis-Maury, Vinay Devadas, and Kesari Mishra, *NetApp*

Pay Migration Tax to Homeland: Anchor-based Scalable Reference Counting for Multicores79
Seokyoung Jung, Jongbin Kim, Minsoo Ryu, Sooyong Kang, and Hyungsoo Jung, *Hanyang University*

Speculative Encryption on GPU Applied to Cryptographic File Systems93
Vandeir Eduardo, *Federal University of Paraná and University of Blumenau*; Luis C. Erpen de Bona and Wagner M. Nunan Zola, *Federal University of Paraná*

Deduplication

Sketching Volume Capacities in Deduplicated Storage107
Danny Harnik and Moshik Hershcovitch, *IBM Research*; Yosef Shatsky, *IBM Systems*; Amir Epstein, *Citi Innovation Lab TLV*; Ronen Kat, *IBM Research*

Finesse: Fine-Grained Feature Locality based Fast Resemblance Detection for Post-Deduplication Delta Compression121
Yucheng Zhang, *Hubei University of Technology*; Wen Xia, *Harbin Institute of Technology, Shenzhen & Peng Cheng Laboratory*; Dan Feng, *WNLO, School of Computer, Huazhong University of Science and Technology*; Hong Jiang, *University of Texas at Arlington*; Yu Hua and Qiang Wang, *WNLO, School of Computer, Huazhong University of Science and Technology*

Sliding Look-Back Window Assisted Data Chunk Rewriting for Improving Deduplication Restore Performance ..129
Zhichao Cao, *University of Minnesota*; Shiyong Liu, *Ocean University of China*; Fenggang Wu, *University of Minnesota*;
Guohua Wang, *South China University of Technology*; Bingzhe Li and David H.C. Du, *University of Minnesota*

Storage Potpourri

DistCache: Provable Load Balancing for Large-Scale Storage Systems with Distributed Caching 143
Zaoxing Liu and Zhihao Bai, *Johns Hopkins University*; Zhenming Liu, *College of William and Mary*; Xiaozhou Li, *Celer Network*; Changhoon Kim, *Barefoot Networks*; Vladimir Braverman and Xin Jin, *Johns Hopkins University*; Ion Stoica, *UC Berkeley*

GearDB: A GC-free Key-Value Store on HM-SMR Drives with Gear Compaction 159
Ting Yao, *Huazhong University of Science and Technology and Temple University*; Jiguang Wan, *Huazhong University of Science and Technology*; Ping Huang, *Temple University*; Yiwen Zhang, Zhiwen Liu, and Changsheng Xie, *Huazhong University of Science and Technology*; Xubin He, *Temple University*

SPEICHER: Securing LSM-based Key-Value Stores using Shielded Execution 173
Maurice Bailleu, Jörg Thalheim, and Pramod Bhatotia, *The University of Edinburgh*; Christof Fetzer, *TU Dresden*; Michio Honda, *NEC Labs*; Kapil Vaswani, *Microsoft Research*

NVM File and Storage Systems

SLM-DB: Single-Level Key-Value Store with Persistent Memory 191
Olzhas Kaiyrakhmet and Songyi Lee, *UNIST*; Beomseok Nam, *Sungkyunkwan University*; Sam H. Noh and Young-ri Choi, *UNIST*

Ziggurat: A Tiered File System for Non-Volatile Main Memories and Disks 207
Shengan Zheng, *Shanghai Jiao Tong University*; Morteza Hoseinzadeh and Steven Swanson, *University of California, San Diego*

Orion: A Distributed File System for Non-Volatile Main Memory and RDMA-Capable Networks 221
Jian Yang, Joseph Izraelevitz, and Steven Swanson, *UC San Diego*

Big Systems

INSTalytics: Cluster Filesystem Co-design for Big-data Analytics 235
Muthian Sivathanu, Midhul Vuppalapati, Bhargav Gulavani, Kaushik Rajan, and Jyoti Leeka, *Microsoft Research India*; Jayashree Mohan, *Univ. of Texas Austin*; Piyus Kedia, *IIT Delhi*

GRAPHONE: A Data Store for Real-time Analytics on Evolving Graphs 249
Pradeep Kumar and H. Howie Huang, *George Washington University*

Automatic, Application-Aware I/O Forwarding Resource Allocation 265
Xu Ji, *Tsinghua University*; *National Supercomputing Center in Wuxi*; Bin Yang and Tianyu Zhang, *National Supercomputing Center in Wuxi*; *Shandong University*; Xiaosong Ma, *Qatar Computing Research Institute, HBKU*; Xiupeng Zhu, *National Supercomputing Center in Wuxi*; *Shandong University*; Xiyang Wang, *National Supercomputing Center in Wuxi*; Nosayba El-Sayed, *Emory University*; Jidong Zhai, *Tsinghua University*; Weiguo Liu, *National Supercomputing Center in Wuxi*; *Shandong University*; Wei Xue, *Tsinghua University*; *National Supercomputing Center in Wuxi*

Flash and Emerging Storage Systems

Design Tradeoffs for SSD Reliability 281
Bryan S. Kim, *Seoul National University*; Jongmoo Choi, *Dankook University*; Sang Lyul Min, *Seoul National University*

Fully Automatic Stream Management for Multi-Streamed SSDs Using Program Contexts 295
Taejin Kim and Duwon Hong, *Seoul National University*; Sangwook Shane Hahn, *Western Digital*; Myoungjun Chun, *Seoul National University*; Sungjin Lee, *DGIST*; Jooyoung Hwang and Jongyoul Lee, *Samsung Electronics*; Jihong Kim, *Seoul National University*

Large-Scale Graph Processing on Emerging Storage Devices 309
Nima Elyasi, *The Pennsylvania State University*; Changho Choi, *Samsung Semiconductor Inc.*; Anand Sivasubramaniam, *The Pennsylvania State University*

(continued on next page)

Erasure Coding and Reliability

Fast Erasure Coding for Data Storage: A Comprehensive Study of the Acceleration Techniques..... 317

Tianli Zhou and Chao Tian, *Texas A&M University*

OpenEC: Toward Unified and Configurable Erasure Coding Management in Distributed Storage Systems..... 331

Xiaolu Li, Runhui Li, and Patrick P. C. Lee, *The Chinese University of Hong Kong*; Yuchong Hu, *Huazhong University of Science and Technology*

Cluster storage systems gotta have HeART: improving storage efficiency by exploiting disk-reliability heterogeneity..... 345

Saurabh Kadekodi, K. V. Rashmi, and Gregory R. Ganger, *Carnegie Mellon University*

ScaleCheck: A Single-Machine Approach for Discovering Scalability Bugs in Large Distributed Systems..... 359

Cesar A. Stuardo, *University of Chicago*; Tanakorn Leesatapornwongsa, *Samsung Research America*; Riza O. Suminto, Huan Ke, and Jeffrey F. Lukman, *University of Chicago*; Wei-Chiu Chuang, *Cloudera*; Shan Lu and Haryadi S. Gunawi, *University of Chicago*