

20th USENIX Conference on File and Storage Technologies (FAST '22)

February 22–24, 2022

Santa Clara, CA, USA

Tuesday, February 22

Persistent Memory: Making it Stick

NyxCache: Flexible and Efficient Multi-tenant Persistent Memory Caching 1
Kan Wu, Kaiwei Tu, and Yuvraj Patel, *University of Wisconsin–Madison*; Rathijit Sen and Kwanghyun Park, *Microsoft*;
Andrea Arpaci-Dusseau and Remzi Arpaci-Dusseau, *University of Wisconsin–Madison*

HTMFMS: Strong Consistency Comes for Free with Hardware Transactional Memory in Persistent Memory File Systems 17
Jifei Yi, Mingkai Dong, Fangnuo Wu, and Haibo Chen, *Institute of Parallel and Distributed Systems, Shanghai Jiao Tong University*

ctFS: Replacing File Indexing with Hardware Memory Translation through Contiguous File Allocation for Persistent Memory 35
Ruibin Li, Xiang Ren, Xu Zhao, Siwei He, Michael Stumm, and Ding Yuan, *University of Toronto*

FORD: Fast One-sided RDMA-based Distributed Transactions for Disaggregated Persistent Memory 51
Ming Zhang, Yu Hua, Pengfei Zuo, and Lurong Liu, *Huazhong University of Science and Technology*

A Series of Merges

Closing the B+-tree vs. LSM-tree Write Amplification Gap on Modern Storage Hardware with Built-in Transparent Compression 69
Yifan Qiao, *Rensselaer Polytechnic Institute*; Xubin Chen, *Google Inc.*; Ning Zheng, Jiangpeng Li, and Yang Liu, *ScaleFlux Inc.*; Tong Zhang, *Rensselaer Polytechnic Institute and ScaleFlux Inc.*

TVStore: Automatically Bounding Time Series Storage via Time-Varying Compression 83
Yanzhe An, *Tsinghua University*; Yue Su, *Huawei Technologies Co., Ltd.*; Yuqing Zhu and Jianmin Wang, *Tsinghua University*

Removing Double-Logging with Passive Data Persistence in LSM-tree based Relational Databases 101
Kecheng Huang, *Shandong University, The Chinese University of Hong Kong*; Zhaoyan Shen and Zhiping Jia, *Shandong University*; Zili Shao, *The Chinese University of Hong Kong*; Feng Chen, *Louisiana State University*

Solidifying the State of SSDs

Improving the Reliability of Next Generation SSDs using WOM-v Codes 117
Shehbaz Jaffer, *University of Toronto, Google*; Kaveh Mahdavian and Bianca Schroeder, *University of Toronto*

GuardedErase: Extending SSD Lifetimes by Protecting Weak Wordlines 133
Duwon Hong, *Seoul National University*; Myungsuk Kim, *Kyungpook National University*; Geonhee Cho, Dusol Lee, and Jihong Kim, *Seoul National University*

Hardware/Software Co-Programmable Framework for Computational SSDs to Accelerate Deep Learning Service on Large-Scale Graphs 147
Miryeong Kwon, Donghyun Gouk, Sangwon Lee, and Myoungsoo Jung, *Computer Architecture and Memory Systems Laboratory, Korea Advanced Institute of Science and Technology (KAIST)*

Operational Characteristics of SSDs in Enterprise Storage Systems: A Large-Scale Field Study 165
Stathis Maneas and Kaveh Mahdavian, *University of Toronto*; Tim Emami, *NetApp*; Bianca Schroeder, *University of Toronto*

Wednesday, February 23

Distant Memories of Efficient Transactions

- Hydra : Resilient and Highly Available Remote Memory** 181
Youngmoon Lee, *Hanyang University*; Hasan Al Maruf and Mosharaf Chowdhury, *University of Michigan*; Asaf Cidon, *Columbia University*; Kang G. Shin, *University of Michigan*
- MT²: Memory Bandwidth Regulation on Hybrid NVM/DRAM Platforms** 199
Jifei Yi, Benchao Dong, Mingkai Dong, Ruizhe Tong, and Haibo Chen, *Institute of Parallel and Distributed Systems, Shanghai Jiao Tong University*
- Aurogon: Taming Aborts in All Phases for Distributed In-Memory Transactions**217
Tianyang Jiang, Guangyan Zhang, Zhiyue Li, and Weimin Zheng, *Tsinghua University*

The Five Ws of Deduplication

- DedupSearch: Two-Phase Deduplication Aware Keyword Search** 233
Nadav Elias, *Technion - Israel Institute of Technology*; Philip Shilane, *Dell Technologies*; Sarai Sheinvald, *ORT Braude College of Engineering*; Gala Yadgar, *Technion - Israel Institute of Technology*
- DeepSketch: A New Machine Learning-Based Reference Search Technique for Post-Deduplication Delta Compression** 247
Jisung Park, *ETH Zürich*; Jeonggyun Kim, Yeseong Kim, and Sungjin Lee, *DGIST*; Onur Mutlu, *ETH Zürich*
- The what, The from, and The to: The Migration Games in Deduplicated Systems** 265
Roei Kisous and Ariel Kolikant, *Technion - Israel Institute of Technology*; Abhinav Duggal, *DELL EMC*; Sarai Sheinvald, *ORT Braude College of Engineering*; Gala Yadgar, *Technion - Israel Institute of Technology*
- DUPEFS: Leaking Data Over the Network With Filesystem Deduplication Side Channels** 281
Andrei Bacs and Saidgani Musaev, *VUsec, Vrije Universiteit Amsterdam*; Kaveh Razavi, *ETH Zurich*; Cristiano Giuffrida and Herbert Bos, *VUsec, Vrije Universiteit Amsterdam*

Meet the 2022 File System Model-Year Lineup

- FusionFS: Fusing I/O Operations using CISC_{Ops} in Firmware File Systems** 297
Jian Zhang, Yujie Ren, and Sudarsun Kannan, *Rutgers University*
- InfiniFS: An Efficient Metadata Service for Large-Scale Distributed Filesystems** 313
Wenhao Lv and Youyou Lu, *Department of Computer Science and Technology, BNRist, Tsinghua University*; Yiming Zhang, *School of Informatics, Xiamen University*; Peile Duan, *Alibaba Group*; Jiwu Shu, *Department of Computer Science and Technology, BNRist, Tsinghua University and School of Informatics, Xiamen University*
- ScaleXFS: Getting scalability of XFS back on the ring** 329
Dohyun Kim, Kwangwon Min, Joontaek Oh, and Youjip Won, *KAIST*
- exF2FS: Transaction Support in Log-Structured Filesystem** 345
Joontaek Oh, Sion Ji, Yongjin Kim, and Youjip Won, *KAIST*

Thursday, February 24

Keys to the Graph Kingdom

- A Log-Structured Merge Tree-aware Message Authentication Scheme for Persistent Key-Value Stores** 363
Igjae Kim, *UNIST, KAIST*; J. Hyun Kim, Minu Chung, Hyungon Moon, and Sam H. Noh, *UNIST*
- Practicably Boosting the Processing Performance of BFS-like Algorithms on Semi-External Graph System via I/O-Efficient Graph Ordering** 381
Tsun-Yu Yang, Yuhong Liang, and Ming-Chang Yang, *The Chinese University of Hong Kong*
- DEPART: Replica Decoupling for Distributed Key-Value Storage** 397
Qiang Zhang and Yongkun Li, *University of Science and Technology of China*; Patrick P. C. Lee, *The Chinese University of Hong Kong*; Yinlong Xu, *Anhui Province Key Laboratory of High Performance Computing, University of Science and Technology of China*; Si Wu, *University of Science and Technology of China*

Keeping the Fast in FAST

PAIO: General, Portable I/O Optimizations With Minor Application Modifications 413

Ricardo Macedo, *INESC TEC and University of Minho*; Yusuke Tanimura and Jason Haga, *AIST*; Vijay Chidambaram, *UT Austin and VMware Research*; José Pereira and João Paulo, *INESC TEC and University of Minho*

Separating Data via Block Invalidation Time Inference for Write Amplification Reduction in Log-Structured Storage 429

Qiuping Wang, *The Chinese University of Hong Kong and Alibaba Group*; Jinhong Li, and Patrick P. C. Lee, *The Chinese University of Hong Kong*; Tao Ouyang, Chao Shi, and Lilong Huang, *Alibaba Group*

CacheSifter: Sifting Cache Files for Boosted Mobile Performance and Lifetime 445

Yu Liang, *Department of Computer Science, City University of Hong Kong and School of Cyber Science and Technology, Zhejiang University*; Riwei Pan, Tianyu Ren, and Yufei Cui, *Department of Computer Science, City University of Hong Kong*; Rachata Ausavarungnirun, *TGGS, King Mongkut's University of Technology North Bangkok*; Xianzhang Chen, *College of Computer Science, Chongqing University*; Changlong Li, *School of Computer Science and Technology, East China Normal University*; Tei-Wei Kuo, *Department of Computer Science, City University of Hong Kong, Department of Computer Science and Information Engineering, National Taiwan University, and NTU High Performance and Scientific Computing Center, National Taiwan University*; Chun Jason Xue, *Department of Computer Science, City University of Hong Kong*