

19th USENIX Conference on File and Storage Technologies (FAST '21)

February 23–25, 2021

Tuesday, February 23

Indexing and Key-Value Store

- ROART: Range-query Optimized Persistent ART** 1
Shaonan Ma and Kang Chen, *Tsinghua University*; Shimin Chen, *SKL of Computer Architecture, ICT, CAS, and University of Chinese Academy of Sciences*; Mengxing Liu, Jianglang Zhu, Hongbo Kang, and Yongwei Wu, *Tsinghua University*
- SpanDB: A Fast, Cost-Effective LSM-tree Based KV Store on Hybrid Storage**17
Hao Chen, *University of Science and Technology of China & Qatar Computing Research Institute, HBKU*; Chaoyi Ruan and Cheng Li, *University of Science and Technology of China*; Xiaosong Ma, *Qatar Computing Research Institute, HBKU*; Yinlong Xu, *University of Science and Technology of China & Anhui Province Key Laboratory of High Performance Computing*
- Evolution of Development Priorities in Key-value Stores Serving Large-scale Applications:
The RocksDB Experience.** 33
Siyang Dong, Andrew Kryczka, and Yanqin Jin, *Facebook Inc.*; Michael Stumm, *University of Toronto*
- REMIX: Efficient Range Query for LSM-trees** 51
Wenshao Zhong, Chen Chen, and Xingbo Wu, *University of Illinois at Chicago*; Song Jiang, *University of Texas at Arlington*

Advanced File Systems

- High Velocity Kernel File Systems with Bento.** 65
Samantha Miller, Kaiyuan Zhang, Mengqi Chen, and Ryan Jennings, *University of Washington*; Ang Chen, *Rice University*; Danyang Zhuo, *Duke University*; Thomas Anderson, *University of Washington*
- Scalable Persistent Memory File System with Kernel-Userspace Collaboration** 81
Youmin Chen, Youyou Lu, and Bohong Zhu, *Tsinghua University*; Andrea C. Arpaci-Dusseau and Remzi H. Arpaci-Dusseau, *University of Wisconsin–Madison*; Jiwu Shu, *Tsinghua University*
- Rethinking File Mapping for Persistent Memory** 97
Ian Neal, Gefei Zuo, Eric Shiple, and Tanvir Ahmed Khan, *University of Michigan*; Youngjin Kwon, *School of Computing, KAIST*; Simon Peter, *University of Texas at Austin*; Baris Kasikci, *University of Michigan*
- pFSCK: Accelerating File System Checking and Repair for Modern Storage.** 113
David Domingo and Sudarsun Kannan, *Rutgers University*
- Pattern-Guided File Compression with User-Experience Enhancement for Log-Structured File System on Mobile Devices.** 127
Cheng Ji, *Nanjing University of Science and Technology*; Li-Pin Chang, *National Chiao Tung University, National Yang Ming Chiao Tung University*; Riwei Pan and Chao Wu, *City University of Hong Kong*; Congming Gao, *Tsinghua University*; Liang Shi, *East China Normal University*; Tei-Wei Kuo and Chun Jason Xue, *City University of Hong Kong*

Wednesday, February 24

Transactions, Deduplication, and More

- ArchTM: Architecture-Aware, High Performance Transaction for Persistent Memory**141
Kai Wu and Jie Ren, *University of California, Merced*; Ivy Peng, *Lawrence Livermore National Laboratory*; Dong Li, *University of California, Merced*
- SPHT: Scalable Persistent Hardware Transactions** 155
Daniel Castro, *INESC-ID & Instituto Superior Técnico*; Alexandro Baldassin, *UNESP - Universidade Estadual Paulista*; João Barreto and Paolo Romano, *INESC-ID & Instituto Superior Técnico*
- The Dilemma between Deduplication and Locality: Can Both be Achieved?**171
Xiangyu Zou and Jingsong Yuan, *Harbin Institute of Technology, Shenzhen*; Philip Shilane, *Dell Technologies*; Wen Xia, *Harbin Institute of Technology, Shenzhen, and Wuhan National Laboratory for Optoelectronics*; Haijun Zhang and Xuan Wang, *Harbin Institute of Technology, Shenzhen*

Remap-SSD: Safely and Efficiently Exploiting SSD Address Remapping to Eliminate Duplicate Writes	187
You Zhou, Qiulin Wu, and Fei Wu, <i>Huazhong University of Science and Technology</i> ; Hong Jiang, <i>University of Texas at Arlington</i> ; Jian Zhou and Changsheng Xie, <i>Huazhong University of Science and Technology</i>	
CheckFreq: Frequent, Fine-Grained DNN Checkpointing	203
Jayashree Mohan, <i>UT Austin</i> ; Amar Phanishayee, <i>Microsoft Research</i> ; Vijay Chidambaram, <i>UT Austin and VMware research</i>	

Cloud and Distributed Systems

Facebook’s Tectonic Filesystem: Efficiency from Exascale	217
Satadru Pan, <i>Facebook, Inc.</i> ; Theano Stavrinos, <i>Facebook, Inc. and Princeton University</i> ; Yunqiao Zhang, Atul Sikaria, Pavel Zakharov, Abhinav Sharma, Shiva Shankar P, Mike Shuey, Richard Wareing, Monika Gangapuram, Guanglei Cao, Christian Preseau, Pratap Singh, Kestutis Patiejunas, and JR Tipton, <i>Facebook, Inc.</i> ; Ethan Katz-Bassett, <i>Columbia University</i> ; Wyatt Lloyd, <i>Princeton University</i>	
Exploiting Combined Locality for Wide-Stripe Erasure Coding in Distributed Storage	233
Yuchong Hu, Liangfeng Cheng, and Qiaori Yao, <i>Huazhong University of Science & Technology</i> ; Patrick P. C. Lee, <i>The Chinese University of Hong Kong</i> ; Weichun Wang and Wei Chen, <i>HIKVISION</i>	
On the Feasibility of Parser-based Log Compression in Large-Scale Cloud Systems	249
Junyu Wei and Guangyan Zhang, <i>Tsinghua University</i> ; Yang Wang, <i>The Ohio State University</i> ; Zhiwei Liu, <i>China University of Geosciences</i> ; Zhanyang Zhu and Junchao Chen, <i>Tsinghua University</i> ; Tingtao Sun and Qi Zhou, <i>Alibaba Cloud</i>	
CNSBench: A Cloud Native Storage Benchmark	263
Alex Merenstein, <i>Stony Brook University</i> ; Vasily Tarasov, Ali Anwar, and Deepavali Bhagwat, <i>IBM Research–Almaden</i> ; Julie Lee, <i>Stony Brook University</i> ; Lukas Rupperecht and Dimitris Skourtis, <i>IBM Research–Almaden</i> ; Yang Yang and Erez Zadok, <i>Stony Brook University</i>	
Concordia: Distributed Shared Memory with In-Network Cache Coherence	277
Qing Wang, Youyou Lu, Erci Xu, Junru Li, Youmin Chen, and Jiwu Shu, <i>Tsinghua University</i>	

Thursday, February 25

Caching Everywhere

eMRC: Efficient Miss Ratio Approximation for Multi-Tier Caching	293
Zhang Liu, <i>University of Colorado Boulder</i> ; Hee Won Lee, <i>Samsung Electronics</i> ; Yu Xiang, <i>AT&T Labs Research</i> ; Dirk Grunwald and Sangtae Ha, <i>University of Colorado Boulder</i>	
The Storage Hierarchy is Not a Hierarchy: Optimizing Caching on Modern Storage Devices with Orthus	307
Kan Wu, Zhihan Guo, Guanzhou Hu, and Kaiwei Tu, <i>University of Wisconsin–Madison</i> ; Ramnathan Alagappan, <i>VMware Research</i> ; Rathijit Sen and Kwanghyun Park, <i>Microsoft</i> ; Andrea C. Arpaci-Dusseau and Remzi H. Arpaci-Dusseau, <i>University of Wisconsin–Madison</i>	
A Community Cache with Complete Information	325
Mania Abdi, <i>Northeastern University</i> ; Amin Mosayyebzadeh, <i>Boston University</i> ; Mohammad Hossein Hajkazemi, <i>Northeastern University</i> ; Emine Ugur Kaynar, <i>Boston University</i> ; Ata Turk, <i>State Street</i> ; Larry Rudolph, <i>TwoSigma</i> ; Orran Krieger, <i>Boston University</i> ; Peter Desnoyers, <i>Northeastern University</i>	
Learning Cache Replacement with CACHEUS	341
Liana V. Rodriguez, Farzana Yusuf, Steven Lyons, Eysler Paz, Raju Rangaswami, and Jason Liu, <i>Florida International University</i> ; Ming Zhao, <i>Arizona State University</i> ; Giri Narasimhan, <i>Florida International University</i>	

The SSD Revolution Is Not Over

FusionRAID: Achieving Consistent Low Latency for Commodity SSD Arrays	355
Tianyang Jiang, Guangyan Zhang, and Zican Huang, <i>Tsinghua University</i> ; Xiaosong Ma, <i>Qatar Computing Research Institute, HBKU</i> ; Junyu Wei, Zhiyue Li, and Weimin Zheng, <i>Tsinghua University</i>	
Behemoth: A Flash-centric Training Accelerator for Extreme-scale DNNs	371
Shine Kim, <i>Seoul National University and Samsung Electronics</i> ; Yunho Jin, Gina Sohn, Jonghyun Bae, Tae Jun Ham, and Jae W. Lee, <i>Seoul National University</i>	

FlashNeuron: SSD-Enabled Large-Batch Training of Very Deep Neural Networks.	387
<i>Jonghyun Bae, Seoul National University; Jongsung Lee, Seoul National University and Samsung Electronics;</i>	
<i>Yunho Jin and Sam Son, Seoul National University; Shine Kim, Seoul National University and Samsung Electronics;</i>	
<i>Hakbeom Jang, Samsung Electronics; Tae Jun Ham and Jae W. Lee, Seoul National University</i>	
D2FQ: Device-Direct Fair Queueing for NVMe SSDs	403
<i>Jiwon Woo, Minwoo Ahn, Gyusun Lee, and Jinkyu Jeong, Sungkyunkwan University</i>	
An In-Depth Study of Correlated Failures in Production SSD-Based Data Centers	417
<i>Shujie Han and Patrick P. C. Lee, The Chinese University of Hong Kong; Fan Xu, Yi Liu, Cheng He, and Jiongzhou Liu,</i>	
<i>Alibaba Group</i>	