

**14th USENIX Conference
on File and Storage Technologies (FAST '16)
February 22–25, 2016
Santa Clara, CA, USA**

Message from the Program Co-Chairs. vii

Tuesday, February 23, 2016

The Blueprint: File and Storage System Designs

Optimizing Every Operation in a Write-optimized File System1

Jun Yuan, Yang Zhan, William Jannen, Prashant Pandey, Amogh Akshintala, Kanchan Chandnani, and Pooja Deo, *Stony Brook University*; Zardosht Kasheff, *Facebook*; Leif Walsh, *Two Sigma*; Michael A. Bender, *Stony Brook University*; Martin Farach-Colton, *Rutgers University*; Rob Johnson, *Stony Brook University*; Bradley C. Kuzmaul, *Massachusetts Institute of Technology*; Donald E. Porter, *Stony Brook University*

The Composite-file File System: Decoupling the One-to-One Mapping of Files and Metadata for Better Performance.15

Shuanglong Zhang, Helen Catanese, and An-I Andy Wang, *Florida State University*

Isotope: Transactional Isolation for Block Storage23

Ji-Yong Shin, *Cornell University*; Mahesh Balakrishnan, *Yale University*; Tudor Marian, *Google*; Hakim Weatherspoon, *Cornell University*

BTrDB: Optimizing Storage System Design for Timeseries Processing39

Michael P Andersen and David E. Culler, *University of California, Berkeley*

Emotional Rescue: Reliability

Environmental Conditions and Disk Reliability in Free-cooled Datacenters.53

Ioannis Manousakis, *Rutgers University*; Sriram Sankar, *GoDaddy*; Gregg McKnight, *Microsoft*; Thu D. Nguyen, *Rutgers University*; Ricardo Bianchini, *Microsoft*

Flash Reliability in Production: The Expected and the Unexpected.67

Bianca Schroeder, *University of Toronto*; Raghav Lagisetty and Arif Merchant, *Google Inc.*

Opening the Chrysalis: On the Real Repair Performance of MSR Codes81

Lluís Pamies-Juarez, Filip Blagojević, Robert Mateescu, and Cyril Gyuot, *WD Research*; Eyal En Gad, *University of Southern California*; Zvonimir Bandic, *WD Research*

They Said It Couldn't Be Done: Writing to Flash

The Devil Is in the Details: Implementing Flash Page Reuse with WOM Codes95

Fabio Margaglia, *Johannes Gutenberg—Universität Mainz*; Gala Yadgar and Eitan Yaakobi, *Technion—Israel Institute of Technology*; Yue Li, *California Institute of Technology*; Assaf Schuster, *Technion—Israel Institute of Technology*; André Brinkmann, *Johannes Gutenberg—Universität Mainz*

Reducing Solid-State Storage Device Write Stress through Opportunistic In-place Delta Compression111

Xuebin Zhang, Jiangpeng Li, and Hao Wang, *Rensselaer Polytechnic Institute*; Kai Zhao, *SanDisk Corporation*; Tong Zhang, *Rensselaer Polytechnic Institute*

Access Characteristic Guided Read and Write Cost Regulation for Performance Improvement on Flash Memory125

Qiao Li and Liang Shi, *Chongqing University*; Chun Jason Xue, *City University of Hong Kong*; Kaijie Wu, *Chongqing University*; Cheng Ji, *City University of Hong Kong*; Qingfeng Zhuge and Edwin H.-M. Sha, *Chongqing University*

Wednesday, February 24, 2016

Songs in the Key of Life: Key-Value Stores

WiscKey: Separating Keys from Values in SSD-conscious Storage.133

Lanyue Lu, Thanumalayan Sankaranarayana Pillai, Andrea C. Arpaci-Dusseau, and Remzi H. Arpaci-Dusseau, *University of Wisconsin—Madison*

Towards Accurate and Fast Evaluation of Multi-Stage Log-structured Designs.149

Hyeontaek Lim and David G. Andersen, *Carnegie Mellon University*; Michael Kaminsky, *Intel Labs*

Efficient and Available In-memory KV-Store with Hybrid Erasure Coding and Replication167

Heng Zhang, Mingkai Dong, and Haibo Chen, *Shanghai Jiao Tong University*

Master of Puppets: Adapting Cloud and Datacenter Storage

Slacker: Fast Distribution with Lazy Docker Containers.181

Tyler Harter, *University of Wisconsin—Madison*; Brandon Salmon and Rose Liu, *Tintri*; Andrea C. Arpaci-Dusseau and Remzi H. Arpaci-Dusseau, *University of Wisconsin—Madison*

sRoute: Treating the Storage Stack Like a Network197

Ioan Stefanovici and Bianca Schroeder, *University of Toronto*; Greg O’Shea, *Microsoft Research*; Eno Thereska, *Confluent and Imperial College London*

Flamingo: Enabling Evolvable HDD-based Near-Line Storage.213

Sergey Legtchenko, Xiaozhou Li, Antony Rowstron, Austin Donnelly, and Richard Black, *Microsoft Research*

Magical Mystery Tour: Miscellaneous

PCAP: Performance-aware Power Capping for the Disk Drive in the Cloud.227

Mohammed G. Khatib and Zvonimir Bandic, *WDC Research*

Mitigating Sync Amplification for Copy-on-write Virtual Disk241

Qingshu Chen, Liang Liang, Yubin Xia, and Haibo Chen, *Shanghai Jiao Tong University*; Hyunsoo Kim, *Samsung Electronics*

Uncovering Bugs in Distributed Storage Systems during Testing (Not in Production!)249

Pantazis Deligiannis, *Imperial College London*; Matt McCutchen, *Massachusetts Institute of Technology*; Paul Thomson, *Imperial College London*; Shuo Chen, *Microsoft*; Alastair F. Donaldson, *Imperial College London*; John Erickson, Cheng Huang, Akash Lal, Rashmi Mudduluru, Shaz Qadeer, and Wolfram Schulte, *Microsoft*

The Tail at Store: A Revelation from Millions of Hours of Disk and SSD Deployments.263

Mingzhe Hao, *University of Chicago*; Gokul Soundararajan and Deepak Kenchammana-Hosekote, *NetApp, Inc.*; Andrew A. Chien and Haryadi S. Gunawi, *University of Chicago*

Thursday, February 25, 2016

Eliminator: Deduplication

Estimating Unseen Deduplication—from Theory to Practice277
Danny Harnik, Ety Khaitzin, and Dmitry Sotnikov, *IBM Research—Haifa*

OrderMergeDedup: Efficient, Failure-Consistent Deduplication on Flash291
Zhuan Chen and Kai Shen, *University of Rochester*

CacheDedup: In-line Deduplication for Flash Caching301
Wenji Li, *Arizona State University*; Gregory Jean-Baptiste, Juan Riveros, and Giri Narasimhan, *Florida International University*; Tony Zhang, *Rensselaer Polytechnic Institute*; Ming Zhao, *Arizona State University*

Using Hints to Improve Inline Block-layer Deduplication315
Sonam Mandal, *Stony Brook University*; Geoff Kuenning, *Harvey Mudd College*; Dongju Ok and Varun Shastri, *Stony Brook University*; Philip Shilane, *EMC Corporation*; Sun Zhen, *Stony Brook University and National University of Defense Technology*; Vasily Tarasov, *IBM Research—Almaden*; Erez Zadok, *Stony Brook University*

The Unforgettable Fire: Flash and NVM

NOVA: A Log-structured File System for Hybrid Volatile/Non-volatile Main Memories323
Jian Xu and Steven Swanson, *University of California, San Diego*

Application-Managed Flash339
Sungjin Lee, Ming Liu, Sangwoo Jun, and Shuotao Xu, *MIT CSAIL*; Jihong Kim, *Seoul National University*; Arvind, *MIT CSAIL*

CloudCache: On-demand Flash Cache Management for Cloud Computing355
Dulcardo Arteaga and Jorge Cabrera, *Florida International University*; Jing Xu, *VMware Inc.*; Swaminathan Sundararaman, *Parallel Machines*; Ming Zhao, *Arizona State University*