

12th USENIX Conference on File and Storage Technologies
February 17–20, 2014
Santa Clara, CA

Message from the Program Co-Chairs. vi

Tuesday, February 18, 2014

Big Memory

Log-structured Memory for DRAM-based Storage1
Stephen M. Rumble, Ankita Kejriwal, and John Ousterhout, *Stanford University*

Strata: High-Performance Scalable Storage on Virtualized Non-volatile Memory17
Brendan Cully, Jake Wires, Dutch Meyer, Kevin Jamieson, Keir Fraser, Tim Deegan, Daniel Stodden,
Geoffrey Lefebvre, Daniel Ferstay, and Andrew Warfield, *Coho Data*

**Evaluating Phase Change Memory for Enterprise Storage Systems:
A Study of Caching and Tiering Approaches**33
Hyojun Kim, Sangeetha Seshadri, Clement L. Dickey, and Lawrence Chiu, *IBM Almaden Research Center*

Flash and SSDs

Wear Unleveling: Improving NAND Flash Lifetime by Balancing Page Endurance47
Xavier Jimenez, David Novo, and Paolo Ienne, *Ecole Polytechnique Fédérale de Lausanne (EPFL)*

**Lifetime Improvement of NAND Flash-based Storage Systems Using Dynamic Program
and Erase Scaling**61
Jaeyong Jeong and Sangwook Shane Hahn, *Seoul National University*; Sungjin Lee, *MIT/CSAIL*; Jihong Kim,
Seoul National University

ReconFS: A Reconstructable File System on Flash Storage75
Youyou Lu, Jiwu Shu, and Wei Wang, *Tsinghua University*

Personal and Mobile

Toward Strong, Usable Access Control for Shared Distributed Data89
Michelle L. Mazurek, Yuan Liang, William Melicher, Manya Sleeper, Lujo Bauer, Gregory R. Ganger, and
Nitin Gupta, *Carnegie Mellon University*; Michael K. Reiter, *University of North Carolina at Chapel Hill*

On the Energy Overhead of Mobile Storage Systems105
Jing Li, *University of California, San Diego*; Anirudh Badam and Ranveer Chandra, *Microsoft Research*;
Steven Swanson, *University of California, San Diego*; Bruce Worthington and Qi Zhang, *Microsoft*

ViewBox: Integrating Local File Systems with Cloud Storage Services119
Yupu Zhang, *University of Wisconsin—Madison*; Chris Dragga, *University of Wisconsin—Madison and
NetApp, Inc.*; Andrea C. Arpaci-Dusseau and Remzi H. Arpaci-Dusseau, *University of Wisconsin—Madison*

(Tuesday, February 18, continues on p. iv)

RAID and Erasure Codes

- CRAID: Online RAID Upgrades Using Dynamic Hot Data Reorganization**133
Alberto Miranda, *Barcelona Supercomputing Center (BSC-CNS)*; Toni Cortes, *Barcelona Supercomputing Center (BSC-CNS) and Technical University of Catalonia (UPC)*
- STAIR Codes: A General Family of Erasure Codes for Tolerating Device and Sector Failures in Practical Storage Systems**147
Mingqiang Li and Patrick P. C. Lee, *The Chinese University of Hong Kong*
- Parity Logging with Reserved Space: Towards Efficient Updates and Recovery in Erasure-coded Clustered Storage**163
Jeremy C. W. Chan, Qian Ding, Patrick P. C. Lee, and Helen H. W. Chan, *The Chinese University of Hong Kong*

Wednesday, February 19, 2014

Experience from Real Systems

- (Big)Data in a Virtualized World: Volume, Velocity, and Variety in Enterprise Datacenters**177
Robert Birke, Mathias Bjoerkqvist, and Lydia Y. Chen, *IBM Research Zurich Lab*; Evgenia Smirni, *College of William and Mary*; Ton Engbersen *IBM Research Zurich Lab*
- From Research to Practice: Experiences Engineering a Production Metadata Database for a Scale Out File System**191
Charles Johnson, Kimberly Keeton, and Charles B. Morrey III, *HP Labs*; Craig A. N. Soules, *Natero*; Alistair Veitch, *Google*; Stephen Bacon, Oskar Batuner, Marcelo Condotta, Hamilton Coutinho, Patrick J. Doyle, Rafael Eichelberger, Hugo Kiehl, Guilherme Magalhaes, James McEvoy, Padmanabhan Nagarajan, Patrick Osborne, Joaquim Souza, Andy Sparkes, Mike Spitzer, Sebastien Tandel, Lincoln Thomas, and Sebastian Zangaro, *HP Storage*
- Analysis of HDFS Under HBase: A Facebook Messages Case Study**199
Tyler Harter, *University of Wisconsin—Madison*; Dhruva Borthakur, Siying Dong, Amitanand Aiyer, and Liyin Tang, *Facebook Inc.*; Andrea C. Arpaci-Dusseau and Remzi H. Arpaci-Dusseau, *University of Wisconsin—Madison*
- Automatic Identification of Application I/O Signatures from Noisy Server-Side Traces**213
Yang Liu, *North Carolina State University*; Raghul Gunasekaran, *Oak Ridge National Laboratory*; Xiaosong Ma, *Qatar Computing Research Institute and North Carolina State University*; Sudharshan S. Vazhkudai, *Oak Ridge National Laboratory*
- Performance and Efficiency**
- Balancing Fairness and Efficiency in Tiered Storage Systems with Bottleneck-Aware Allocation**229
Hui Wang and Peter Varman, *Rice University*
- SpringFS: Bridging Agility and Performance in Elastic Distributed Storage**243
Lianghong Xu, James Cipar, Elie Krevat, Alexey Tumanov, and Nitin Gupta, *Carnegie Mellon University*; Michael A. Kozuch, *Intel Labs*; Gregory R. Ganger, *Carnegie Mellon University*
- Migratory Compression: Coarse-grained Data Reordering to Improve Compressibility**257
Xing Lin, *University of Utah*; Guanlin Lu, Fred Douglass, Philip Shilane, and Grant Wallace, *EMC Corporation—Data Protection and Availability Division*

Thursday, February 20, 2014

OS and Storage Interactions

Resolving Journaling of Journal Anomaly in Android I/O: Multi-Version B-tree with Lazy Split273

Wook-Hee Kim and Beomseok Nam, *Ulsan National Institute of Science and Technology*; Dongil Park and Youjip Won, *Hanyang University*

Journaling of Journal Is (Almost) Free287

Kai Shen, Stan Park, and Meng Zhu, *University of Rochester*

Checking the Integrity of Transactional Mechanisms295

Daniel Fryer, Mike Qin, Jack Sun, Kah Wai Lee, Angela Demke Brown, and Ashvin Goel, *University of Toronto*

OS and Peripherals

DC Express: Shortest Latency Protocol for Reading Phase Change Memory over PCI Express309

Dejan Vučinić, Qingbo Wang, Cyril Guyot, Robert Mateescu, Filip Blagojević, Luiz Franca-Neto, and Damien Le Moal, *HGST San Jose Research Center*; Trevor Bunker, Jian Xu, and Steven Swanson, *University of California, San Diego*; Zvonimir Bandić, *HGST San Jose Research Center*

MultiLanes: Providing Virtualized Storage for OS-level Virtualization on Many Cores317

Junbin Kang, Benlong Zhang, Tianyu Wo, Chunming Hu, and Jinpeng Huai, *Beihang University*