

**USENIX ATC '13:
2013 USENIX Annual Technical Conference
June 26–28, 2013
San Jose, CA**

Message from the Program Co-Chairs. vi

Wednesday, June 26, 2013

Virtual Machine Implementation

Optimizing VM Checkpointing for Restore Performance in VMware ESXi1
Irene Zhang, *University of Washington and VMware*; Tyler Denniston, *MIT CSAIL and VMware*; Yury Baskakov, *VMware*; Alex Garthwaite, *CloudPhysics and VMware*

Hyper-Switch: A Scalable Software Virtual Switching Architecture13
Kaushik Kumar Ram, Alan L. Cox, Mehul Chadha, and Scott Rixner, *Rice University*

MiG: Efficient Migration of Desktop VMs Using Semantic Compression25
Anshul Rai and Ram Ramjee, *Microsoft Research India*; Ashok Anand, *Bell Labs India*; Venkata N. Padmanabhan, *Microsoft Research India*; George Varghese, *Microsoft Research US*

Computing in the Cloud

Copysets: Reducing the Frequency of Data Loss in Cloud Storage37
Asaf Cidon, Stephen Rumble, Ryan Stutsman, Sachin Katti, John Ousterhout, and Mendel Rosenblum, *Stanford University*

TAO: Facebook's Distributed Data Store for the Social Graph49
Nathan Bronson, Zach Amsden, George Cabrera, Prasad Chakka, Peter Dimov, Hui Ding, Jack Ferris, Anthony Giardullo, Sachin Kulkarni, Harry Li, Mark Marchukov, Dmitri Petrov, Lovro Puzar, Yee Jiun Song, and Venkat Venkataramani, *Facebook, Inc.*

PIKACHU: How to Rebalance Load in Optimizing MapReduce On Heterogeneous Clusters61
Rohan Gandhi, Di Xie, and Y. Charlie Hu, *Purdue University*

Flash-based Storage

FlashFQ: A Fair Queuing I/O Scheduler for Flash-Based SSDs67
Kai Shen and Stan Park, *University of Rochester*

The Harey Tortoise: Managing Heterogeneous Write Performance in SSDs79
Laura M. Grupp, *University of California, San Diego*; John D. Davis, *Microsoft Research*; Steven Swanson, *University of California, San Diego*

Janus: Optimal Flash Provisioning for Cloud Storage Workloads91
Christoph Albrecht, Arif Merchant, Murray Stokely, Muhammad Waliji, François Labelle, Nate Coehlo, Xudong Shi, and C. Eric Schrock, *Google, Inc.*

(Wednesday, June 26, continues on p. iv)

Miscellanea #1

Using One-Sided RDMA Reads to Build a Fast, CPU-Efficient Key-Value Store	103
Christopher Mitchell, <i>New York University</i> ; Yifeng Geng, <i>Tsinghua University</i> ; Jinyang Li, <i>New York University</i>	
Lightweight Memory Tracing	115
Mathias Payer, Enrico Kravina, and Thomas R. Gross, <i>ETH Zurich</i>	
Flash Caching on the Storage Client	127
David A. Holland, Elaine Angelino, Gideon Wald, and Margo I. Seltzer, <i>Harvard University</i>	
Practical and Effective Sandboxing for Non-root Users	139
Taesoo Kim and Nickolai Zeldovich, <i>MIT CSAIL</i>	

Thursday, June 27, 2013

Data Storage

TABLEFS: Enhancing Metadata Efficiency in the Local File System	145
Kai Ren and Garth Gibson, <i>Carnegie Mellon University</i>	
Characterization of Incremental Data Changes for Efficient Data Protection	157
Hyong Shim, Philip Shilane, and Windsor Hsu, <i>EMC Corporation</i>	
On the Efficiency of Durable State Machine Replication	169
Alysson Bessani, Marcel Santos, João Felix, and Nuno Neves, <i>FCUL/LaSIGE, University of Lisbon</i> ; Miguel Correia, <i>INESC-ID, IST, University of Lisbon</i>	
Estimating Duplication by Content-based Sampling	181
Fei Xie, Michael Condict, and Sandip Shete, <i>NetApp Inc.</i>	

Miscellanea #2

MutantX-S: Scalable Malware Clustering Based on Static Features	187
Xin Hu, <i>IBM T.J. Watson Research Center</i> ; Sandeep Bhatkar and Kent Griffin, <i>Symantec Research Labs</i> ; Kang G. Shin, <i>University of Michigan</i>	
Redundant State Detection for Dynamic Symbolic Execution	199
Suhabe Bugarra and Dawson Engler, <i>Stanford University</i>	
packetdrill: Scriptable Network Stack Testing, from Sockets to Packets	213
Neal Cardwell, Yuchung Cheng, Lawrence Brakmo, Matt Mathis, Barath Raghavan, Nandita Dukkupati, Hsiao-keng Jerry Chu, Andreas Terzis, and Tom Herbert, <i>Google</i>	

Virtual Machine Performance

DeepDive: Transparently Identifying and Managing Performance Interference in Virtualized Environments	219
Dejan Novaković, Nedeljko Vasić, and Stanko Novaković, <i>École Polytechnique Fédérale de Lausanne (EPFL)</i> ; Dejan Kostić, <i>Institute IMDEA Networks</i> ; Ricardo Bianchini, <i>Rutgers University</i>	
Efficient and Scalable Paravirtual I/O System	231
Nadav Har'El, Abel Gordon, and Alex Landau, <i>IBM Research-Haifa</i> ; Muli Ben-Yehuda, <i>Technion IIT and Hypervisor Consulting</i> ; Avishay Traeger and Razya Ladelsky, <i>IBM Research-Haifa</i>	
vTurbo: Accelerating Virtual Machine I/O Processing Using Designated Turbo-Sliced Core	243
Cong Xu, Sahan Gamage, Hui Lu, Ramana Kompella, and Dongyan Xu, <i>Purdue University</i>	

Managing Resources

- When Slower Is Faster: On Heterogeneous Multicores for Reliable Systems**255
Tomas Hruby, Herbert Bos, and Andrew S. Tanenbaum, *VU University Amsterdam*
- IAMEM: Interaction-Aware Memory Energy Management**267
Mingsong Bi, *Intel Corporation*; Srinivasan Chandrasekharan, and Chris Gniady, *University of Arizona*
- XLH: More Effective Memory Deduplication Scanners Through Cross-layer Hints**279
Konrad Miller, Fabian Franz, Marc Rittinghaus, Marius Hillenbrand, and Frank Bellosa, *Karlsruhe Institute of Technology*
- Enabling OS Research by Inferring Interactions in the Black-Box GPU Stack**291
Konstantinos Menychtas, Kai Shen, and Michael L. Scott, *University of Rochester*

Friday, June 28, 2013

Small Applications

- Mantis: Automatic Performance Prediction for Smartphone Applications**297
Yongin Kwon, *Seoul National University*; Sangmin Lee, *University of Texas at Austin*; Hayoon Yi, Donghyun Kwon, and Seungjun Yang, *Seoul National University*; Byung-Gon Chun, *Microsoft*; Ling Huang and Petros Maniatis, *Intel*; Mayur Naik, *Georgia Institute of Technology*; Yunheung Paek, *Seoul National University*
- IO Stack Optimization for Smartphones**309
Sooman Jeong, *Hanyang University*; Kisung Lee, *Samsung Electronics*; Seongjin Lee, *Hanyang University*; Seoungbum Son, *Samsung Electronics*; Youjip Won, *Hanyang University*
- How to Run POSIX Apps in a Minimal Picoprocess**321
Jon Howell, Bryan Parno, and John R. Douceur, *Microsoft Research*

Packets

- Network Interface Design for Low Latency Request-Response Protocols**333
Mario Flajslik and Mendel Rosenblum, *Stanford University*
- DEFINED: Deterministic Execution for Interactive Control-Plane Debugging**347
Chia-Chi Lin, Virajith Jalaparti, and Matthew Caesar, *University of Illinois at Urbana-Champaign*; Jacobus Van der Merwe, *University of Utah*
- Improving Server Application Performance via Pure TCP ACK Receive Optimization**359
Michael Chan and David R. Cheriton, *Stanford University*