

11th USENIX Symposium on Operating Systems Design and Implementation (OSDI'14)

Broomfield, Colorado, USA
6 - 8 October 2014

ISBN: 978-1-7138-0465-9

Printed from e-media with permission by:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571



Some format issues inherent in the e-media version may also appear in this print version.

Copyright© (2014) by Usenix Association
All rights reserved.

Printed with permission by Curran Associates, Inc. (2020)

For permission requests, please contact Usenix Association
at the address below.

Usenix Association
2560 Ninth Street, Suite 215
Berkeley, California, 94710

<https://www.usenix.org/>

Additional copies of this publication are available from:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: 845-758-0400
Fax: 845-758-2633
Email: curran@proceedings.com
Web: www.proceedings.com

**11th USENIX Symposium on
Operating Systems Design and Implementation
October 6–8, 2014
Broomfield, CO**

Message from the Program Chair ix

Monday, October 6, 2014

Who Put the Kernel in My OS Conference?

Arrakis: The Operating System is the Control Plane1
Simon Peter, Jialin Li, Irene Zhang, Dan R. K. Ports, Doug Woos, Arvind Krishnamurthy, and Thomas Anderson,
University of Washington; Timothy Roscoe, *ETH Zürich*

Decoupling Cores, Kernels, and Operating Systems17
Gerd Zellweger, Simon Gerber, Kornilios Kourtis, and Timothy Roscoe, *ETH Zürich*

Jitk: A Trustworthy In-Kernel Interpreter Infrastructure33
Xi Wang, David Lazar, Nikolai Zeldovich, and Adam Chlipala, *MIT CSAIL*; Zachary Tatlock, *University
of Washington*

IX: A Protected Dataplane Operating System for High Throughput and Low Latency49
Adam Belay, *Stanford University*; George Prekas, *École Polytechnique Fédérale de Lausanne (EPFL)*;
Ana Klimovic, Samuel Grossman, and Christos Kozyrakis, *Stanford University*; Edouard Bugnion, *École
Polytechnique Fédérale de Lausanne (EPFL)*

Data in the Abstract

Willow: A User-Programmable SSD67
Sudharsan Seshadri, Mark Gahagan, Sundaram Bhaskaran, Trevor Bunker, Arup De, Yanqin Jin, Yang Liu,
and Steven Swanson, *University of California, San Diego*

Physical Disentanglement in a Container-Based File System81
Lanyue Lu, Yupu Zhang, Thanh Do, Samer Al-Kiswany, Andrea C. Arpaci-Dusseau, and Remzi H. Arpaci-
Dusseau, *University of Wisconsin—Madison*

Customizable and Extensible Deployment for Mobile/Cloud Applications97
Irene Zhang, Adriana Szekeres, Dana Van Aken, and Isaac Ackerman, *University of Washington*; Steven D.
Gribble, *Google and University of Washington*; Arvind Krishnamurthy and Henry M. Levy, *University of
Washington*

Pebbles: Fine-Grained Data Management Abstractions for Modern Operating Systems113
Riley Spahn and Jonathan Bell, *Columbia University*; Michael Lee, *The University of Texas at Austin*; Sravan
Bhamidipati, Roxana Geambasu, and Gail Kaiser, *Columbia University*

My Insecurities

Protecting Users by Confining JavaScript with COWL131
Deian Stefan and Edward Z. Yang, *Stanford University*; Petr Marchenko, *Google*; Alejandro Russo, *Chalmers
University of Technology*; Dave Herman, *Mozilla*; Brad Karp, *University College London*; David Mazières,
Stanford University

Code-Pointer Integrity147
Volodymyr Kuznetsov, *École Polytechnique Fédérale de Lausanne (EPFL)*; László Szekeres, *Stony Brook
University*; Mathias Payer, *Purdue University*; George Candea, *École Polytechnique Fédérale de Lausanne
(EPFL)*; R. Sekar, *Stony Brook University*; Dawn Song, *University of California, Berkeley*

(Monday, October 6, continues on next page)

Ironclad Apps: End-to-End Security via Automated Full-System Verification	165
Chris Hawblitzel, Jon Howell, and Jacob R. Lorch, <i>Microsoft Research</i> ; Arjun Narayan, <i>University of Pennsylvania</i> ; Bryan Parno, <i>Microsoft Research</i> ; Danfeng Zhang, <i>Cornell University</i> ; Brian Zill, <i>Microsoft Research</i>	
SHLL: A Secure Shell Scripting Language	183
Scott Moore, Christos Dimoulas, Dan King, and Stephen Chong, <i>Harvard University</i>	

Variety Pack

GPUnet: Networking Abstractions for GPU Programs	201
Sangman Kim, Seonggu Huh, Yige Hu, Xinya Zhang, and Emmett Witchel, <i>The University of Texas at Austin</i> ; Amir Wated and Mark Silberstein, <i>Technion—Israel Institute of Technology</i>	
The Mystery Machine: End-to-end Performance Analysis of Large-scale Internet Services	217
Michael Chow, <i>University of Michigan</i> ; David Meisner, <i>Facebook, Inc.</i> ; Jason Flinn, <i>University of Michigan</i> ; Daniel Peek, <i>Facebook, Inc.</i> ; Thomas F. Wenisch, <i>University of Michigan</i>	
End-to-end Performance Isolation Through Virtual Datacenters	233
Sebastian Angel, <i>The University of Texas at Austin</i> ; Hitesh Ballani, Thomas Karagiannis, Greg O’Shea, and Eno Thereska, <i>Microsoft Research</i>	
Simple Testing Can Prevent Most Critical Failures: An Analysis of Production Failures in Distributed Data-Intensive Systems	249
Ding Yuan, Yu Luo, Xin Zhuang, Guilherme Renna Rodrigues, Xu Zhao, Yongle Zhang, Pranay U. Jain, and Michael Stumm, <i>University of Toronto</i>	

Tuesday, October 7, 2014

Head in the Cloud

Shielding Applications from an Untrusted Cloud with Haven	267
Andrew Baumann, Marcus Peinado, and Galen Hunt, <i>Microsoft Research</i>	
Apollo: Scalable and Coordinated Scheduling for Cloud-Scale Computing	285
Eric Boutin, Jaliya Ekanayake, Wei Lin, Bing Shi, and Jingren Zhou, <i>Microsoft</i> ; Zhengping Qian, Ming Wu, and Lidong Zhou, <i>Microsoft Research</i>	
The Power of Choice in Data-Aware Cluster Scheduling	301
Shivaram Venkataraman and Aurojit Panda, <i>University of California, Berkeley</i> ; Ganesh Ananthanarayanan, <i>Microsoft Research</i> ; Michael J. Franklin and Ion Stoica, <i>University of California, Berkeley</i>	
Heading Off Correlated Failures through Independence-as-a-Service	317
Ennan Zhai, <i>Yale University</i> ; Ruichuan Chen, <i>Bell Labs and Alcatel-Lucent</i> ; David Isaac Wolinsky and Bryan Ford, <i>Yale University</i>	

Storage Runs Hot and Cold

Characterizing Storage Workloads with Counter Stacks	335
Jake Wires, Stephen Ingram, Zachary Drudi, Nicholas J. A. Harvey, and Andrew Warfield, <i>Coho Data</i>	
Pelican: A Building Block for Exascale Cold Data Storage	351
Shobana Balakrishnan, Richard Black, Austin Donnelly, Paul England, Adam Glass, Dave Harper, and Sergey Legtchenko, <i>Microsoft Research</i> ; Aaron Ogus, <i>Microsoft</i> ; Eric Peterson and Antony Rowstron, <i>Microsoft Research</i>	
A Self-Configurable Geo-Replicated Cloud Storage System	367
Masoud Saeida Ardekani, <i>INRIA and Sorbonne Universités</i> ; Douglas B. Terry, <i>Microsoft Research</i>	

f4: Facebook’s Warm BLOB Storage System	383
Subramanian Muralidhar, <i>Facebook, Inc.</i> ; Wyatt Lloyd, <i>University of Southern California and Facebook, Inc.</i> ; Sabyasachi Roy, Cory Hill, Ernest Lin, Weiwen Liu, Satadru Pan, Shiva Shankar, and Viswanath Sivakumar, <i>Facebook, Inc.</i> ; Linpeng Tang, <i>Princeton University and Facebook, Inc.</i> ; Sanjeev Kumar, <i>Facebook, Inc.</i>	

Pest Control

SAMC: Semantic-Aware Model Checking for Fast Discovery of Deep Bugs in Cloud Systems	399
Tanakorn Leesatapornwongsa and Mingzhe Hao, <i>University of Chicago</i> ; Pallavi Joshi, <i>NEC Labs America</i> ; Jeffrey F. Lukman, <i>Surya University</i> ; Haryadi S. Gunawi, <i>University of Chicago</i>	

SKI: Exposing Kernel Concurrency Bugs through Systematic Schedule Exploration	415
Pedro Fonseca, <i>Max Planck Institute for Software Systems (MPI-SWS)</i> ; Rodrigo Rodrigues, <i>CITI/NOVA University of Lisbon</i> ; Björn B. Brandenburg, <i>Max Planck Institute for Software Systems (MPI-SWS)</i>	

All File Systems Are Not Created Equal: On the Complexity of Crafting Crash-Consistent Applications ...	433
Thanumalayan Sankaranarayana Pillai, Vijay Chidambaram, Ramnathan Alagappan, Samer Al-Kiswany, Andrea C. Arpaci-Dusseau, and Remzi H. Arpaci-Dusseau, <i>University of Wisconsin–Madison</i>	

Torturing Databases for Fun and Profit	449
Mai Zheng, <i>The Ohio State University</i> ; Joseph Tucek, <i>HP Labs</i> ; Dachuan Huang and Feng Qin, <i>The Ohio State University</i> ; Mark Lillibridge, Elizabeth S. Yang, and Bill W. Zhao, <i>HP Labs</i> ; Shashank Singh, <i>The Ohio State University</i>	

Transaction Action

Fast Databases with Fast Durability and Recovery Through Multicore Parallelism	465
Wenting Zheng and Stephen Tu, <i>Massachusetts Institute of Technology</i> ; Eddie Kohler, <i>Harvard University</i> ; Barbara Liskov, <i>Massachusetts Institute of Technology</i>	

Extracting More Concurrency from Distributed Transactions	479
Shuai Mu, <i>Tsinghua University and New York University</i> ; Yang Cui and Yang Zhang, <i>New York University</i> ; Wyatt Lloyd, <i>University of Southern California and Facebook, Inc.</i> ; Jinyang Li, <i>New York University</i>	

Salt: Combining ACID and BASE in a Distributed Database	495
Chao Xie, Chunzhi Su, Manos Kapritsos, Yang Wang, Navid Yaghmazadeh, Lorenzo Alvisi, and Prince Mahajan, <i>The University of Texas at Austin</i>	

Phase Reconciliation for Contended In-Memory Transactions	511
Neha Narula and Cody Cutler, <i>MIT CSAIL</i> ; Eddie Kohler, <i>Harvard University</i> ; Robert Morris, <i>MIT CSAIL</i>	

Wednesday, October 8, 2014

Play It Again, Sam

Eidetic Systems	525
David Devecsery, Michael Chow, Xianzheng Dou, Jason Flinn, and Peter M. Chen, <i>University of Michigan</i>	

Detecting Covert Timing Channels with Time-Deterministic Replay	541
Ang Chen, <i>University of Pennsylvania</i> ; W. Brad Moore, <i>Georgetown University</i> ; Hanjun Xiao, Andreas Haeberlen, and Linh Thi Xuan Phan, <i>University of Pennsylvania</i> ; Micah Sherr and Wenchao Zhou, <i>Georgetown University</i>	

Identifying Information Disclosure in Web Applications with Retroactive Auditing	555
Haogang Chen, Taesoo Kim, Xi Wang, Nikolai Zeldovich, and M. Frans Kaashoek, <i>MIT CSAIL</i>	

(Wednesday, October 8, continues on next page)

Help Me Learn

Project Adam: Building an Efficient and Scalable Deep Learning Training System571
Trishul Chilimbi, Yutaka Suzue, Johnson Apacible, and Karthik Kalyanaraman, *Microsoft Research*

Scaling Distributed Machine Learning with the Parameter Server583
Mu Li, *Carnegie Mellon University and Baidu*; David G. Andersen and Jun Woo Park, *Carnegie Mellon University*; Alexander J. Smola, *Carnegie Mellon University and Google, Inc.*; Amr Ahmed, Vanja Josifovski, James Long, Eugene J. Shekita, and Bor-Yiing Su, *Google, Inc.*

GraphX: Graph Processing in a Distributed Dataflow Framework599
Joseph E. Gonzalez, *University of California, Berkeley*; Reynold S. Xin, *University of California, Berkeley, and Databricks*; Ankur Dave, Daniel Crankshaw, and Michael J. Franklin, *University of California, Berkeley*; Ion Stoica, *University of California, Berkeley, and Databricks*

Hammers and Saws

Nail: A Practical Tool for Parsing and Generating Data Formats615
Julian Bangert and Nikolai Zeldovich, *MIT CSAIL*

***lprof*: A Non-intrusive Request Flow Profiler for Distributed Systems**629
Xu Zhao, Yongle Zhang, David Lion, Muhammad Faizan Ullah, Yu Luo, Ding Yuan, and Michael Stumm, *University of Toronto*

Pydron: Semi-Automatic Parallelization for Multi-Core and the Cloud645
Stefan C. Müller, *ETH Zürich and University of Applied Sciences Northwestern Switzerland*; Gustavo Alonso and Adam Amara, *ETH Zürich*; André Csillaghy, *University of Applied Sciences Northwestern Switzerland*

User-Guided Device Driver Synthesis661
Leonid Ryzhyk, *University of Toronto, NICTA, and University of New South Wales*; Adam Walker, *NICTA and University of New South Wales*; John Keys, *Intel Corporation*; Alexander Legg, *NICTA and University of New South Wales*; Arun Raghunath, *Intel Corporation*; Michael Stumm, *University of Toronto*; Mona Vij, *Intel Corporation*