

**USENIX ATC '16**  
**2016 USENIX Annual Technical Conference**  
**June 22–24, 2016**  
**Denver, CO, USA**

Message from the Program Co-Chairs..... vii

**Wednesday, June 22, 2016**

**Datacenter Networking**

**FLICK: Developing and Running Application-Specific Network Services .....1**

Abdul Alim, Richard G. Clegg, Luo Mai, Lukas Rupprecht, and Eric Seckler, *Imperial College London*; Paolo Costa, *Microsoft Research and Imperial College London*; Peter Pietzuch and Alexander L. Wolf, *Imperial College London*; Nik Sultana, Jon Crowcroft, Anil Madhavapeddy, Andrew W. Moore, and Richard Mortier, *University of Cambridge*; Masoud Koleni, Luis Oviedo, and Derek McAuley, *University of Nottingham*; Matteo Migliavacca, *University of Kent*

**SoftFlow: A Middlebox Architecture for Open vSwitch.....15**

Ethan J. Jackson, *University of California, Berkeley*; Melvin Walls, *Penn State Harrisburg and University of California, Berkeley*; Aurojit Panda, *University of California, Berkeley*; Justin Pettit, Ben Pfaff, and Jarno Rajahalme, *VMware, Inc.*; Teemu Koponen, *Styra, Inc.*; Scott Shenker, *University of California, Berkeley, and International Computer Science Institute*

**Fast and Cautious: Leveraging Multi-path Diversity for Transport Loss Recovery in Data Centers .....29**

Guo Chen, *Tsinghua University and Microsoft Research*; Yuanwei Lu, *University of Science and Technology of China and Microsoft Research*; Yuan Meng, *Tsinghua University*; Bojie Li, *University of Science and Technology of China and Microsoft Research*; Kun Tan, *Microsoft Research*; Dan Pei, *Tsinghua University*; Peng Cheng, Layong (Larry) Luo, and Yongqiang Xiong, *Microsoft Research*; Xiaoliang Wang, *Nanjing University*; Youjian Zhao, *Tsinghua University*

**StackMap: Low-Latency Networking with the OS Stack and Dedicated NICs .....43**

Kenichi Yasukata, *Keio University*; Michio Honda, Douglas Santry, and Lars Eggert, *NetApp*

**File and Key-Value Systems**

**SLIK: Scalable Low-Latency Indexes for a Key-Value Store .....57**

Ankita Kejriwal, Arjun Gopalan, Ashish Gupta, Zhihao Jia, Stephen Yang, and John Ousterhout, *Stanford University*

**Understanding Manycore Scalability of File Systems.....71**

Changwoo Min, Sanidhya Kashyap, Steffen Maass, Woonhak Kang, and Taesoo Kim, *Georgia Institute of Technology*

**ParaFS: A Log-Structured File System to Exploit the Internal Parallelism of Flash Devices .....87**

Jiacheng Zhang, Jiwu Shu, and Youyou Lu, *Tsinghua University*

**FastCDC: a Fast and Efficient Content-Defined Chunking Approach for Data Deduplication.....101**

Wen Xia, *Huazhong University of Science and Technology and Sangfor Technologies Co., Ltd.*; Yukun Zhou, *Huazhong University of Science and Technology*; Hong Jiang, *University of Texas at Arlington*; Dan Feng, Yu Hua, Yuchong Hu, Yucheng Zhang, and Qing Liu, *Huazhong University of Science and Technology*

*(Wednesday, June 22 continues on the next page)*

## Mobile and Apps

- Unsafe Time Handling in Smartphones** .....115  
Abhilash Jindal, Prahlad Joshi, Y. Charlie Hu, and Samuel Midkiff, *Purdue University*
- Energy Discounted Computing on Multicore Smartphones** .....129  
Meng Zhu and Kai Shen, *University of Rochester*
- Beam: Ending Monolithic Applications for Connected Devices** .....143  
Chenguang Shen, *University of California, Los Angeles*; Rayman Preet Singh, *Samsung Research*;  
Amar Phanishayee, Aman Kansal, and Ratul Mahajan, *Microsoft Research*
- Caching Doesn't Improve Mobile Web Performance (Much)** .....159  
Jamshed Vesuna and Colin Scott, *University of California, Berkeley*; Michael Buettner and Michael Piatek,  
*Google*; Arvind Krishnamurthy, *University of Washington*; Scott Shenker, *University of California, Berkeley*,  
and *International Computer Science Institute*

## Systems and Network Security

- Secure and Efficient Application Monitoring and Replication** .....167  
Stijn Voleckaert, *University of California, Irvine, and Ghent University*; Bart Coppens, *Ghent University*;  
Alexios Voulimeneas, *University of California, Irvine*; Andrei Homescu, *Immunant, Inc.*; Per Larsen, *University*  
*of California, Irvine, and Immunant, Inc.*; Bjorn De Sutter, *Ghent University*; Michael Franz, *University of*  
*California, Irvine*
- Blockstack: A Global Naming and Storage System Secured by Blockchains** .....181  
Muneeb Ali and Jude Nelson, *Princeton University and Blockstack Labs*; Ryan Shea, *Blockstack Labs*;  
Michael J. Freedman, *Princeton University*
- Satellite: Joint Analysis of CDNs and Network-Level Interference** .....195  
Will Scott, Thomas Anderson, Tadayoshi Kohno, and Arvind Krishnamurthy, *University of Washington*
- Subversive-C: Abusing and Protecting Dynamic Message Dispatch** .....209  
Julian Lettner, *University of California, Irvine*; Benjamin Kollenda, *Ruhr-Universität Bochum*; Andrei Homescu,  
*Immunant, Inc.*; Per Larsen, *University of California, Irvine, and Immunant, Inc.*; Felix Schuster, *Microsoft*  
*Research*; Lucas Davi  
and Ahmad-Reza Sadeghi, *Technische Universität Darmstadt*; Thorsten Holz, *Ruhr-Universität Bochum*;  
Michael Franz, *University of California, Irvine*

## Thursday, June 23, 2016

### Cloud, Coordination, and Consensus

- Callinicos: Robust Transactional Storage for Distributed Data Structures** .....223  
Ricardo Padilha, Enrique Fynn, Robert Soulé, and Fernando Pedone, *Università della Svizzera Italiana (USI)*
- Filo: Consolidated Consensus as a Cloud Service** .....237  
Parisa Jalili Marandi, Christos Gkantsidis, Flavio Junqueira, and Dushyanth Narayanan, *Microsoft Research*
- Modular Composition of Coordination Services** .....251  
Kfir Lev-Ari, *Technion—Israel Institute of Technology*; Edward Bortnikov, *Yahoo Research*; Idit Keidar,  
*Technion—Israel Institute of Technology and Yahoo Research*; Alexander Shraer, *Google*
- Cheap and Available State Machine Replication** .....265  
Rong Shi and Yang Wang, *The Ohio State University*

### Architectural Interaction

- Horton Tables: Fast Hash Tables for In-Memory Data-Intensive Computing** .....281  
Alex D. Breslow, *AMD Research and University of California, San Diego*; Dong Ping Zhang, Joseph L.  
Greathouse, and Nuwan Jayasena, *AMD Research*; Dean M. Tullsen, *University of California, San Diego*

<b>Ginseng: Market-Driven LLC Allocation</b> .....	<b>295</b>
Liran Funaro, Orna Agmon Ben-Yehuda, and Assaf Schuster, <i>Technion—Israel Institute of Technology</i>	
<b>Elfen Scheduling: Fine-Grain Principled Borrowing from Latency-Critical Workloads Using Simultaneous Multithreading</b> .....	<b>309</b>
Xi Yang and Stephen M. Blackburn, <i>Australian National University</i> ; Kathryn S. McKinley, <i>Microsoft Research</i>	
<b>Coherence Stalls or Latency Tolerance: Informed CPU Scheduling for Socket and Core Sharing</b> .....	<b>323</b>
Sharanyan Srikanthan, Sandhya Dwarkadas, and Kai Shen, <i>University of Rochester</i>	
<b>Caching and Indexing</b>	
<b>Replex: A Scalable, Highly Available Multi-Index Data Store</b> .....	<b>337</b>
Amy Tai, <i>VMWare Research and Princeton University</i> ; Michael Wei, <i>VMware Research and University of California, San Diego</i> ; Michael J. Freedman, <i>Princeton University</i> ; Ittai Abraham and Dahlia Malkhi, <i>VMWare Research</i>	
<b>Kinetic Modeling of Data Eviction in Cache</b> .....	<b>351</b>
Xiameng Hu, Xiaolin Wang, Lan Zhou, Yingwei Luo, <i>Peking University</i> ; Chen Ding, <i>University of Rochester</i> ; Zhenlin Wang, <i>Michigan Technological University</i>	
<b>Scalable In-Memory Transaction Processing with HTM</b> .....	<b>365</b>
Yingjun Wu and Kian-Lee Tan, <i>National University of Singapore</i>	
<b>Erasing Belady’s Limitations: In Search of Flash Cache Offline Optimality</b> .....	<b>379</b>
Yue Cheng, <i>Virginia Polytechnic Institute and State University</i> ; Fred Douglass, Philip Shilane, Michael Trachtman, and Grant Wallace, <i>EMC Corporation</i> ; Peter Desnoyers, <i>Northeastern University</i> ; Kai Li, <i>Princeton University</i>	
<b>Energy vs. Performance</b>	
<b>Unlocking Energy</b> .....	<b>393</b>
Babak Falsafi, Rachid Guerraoui, Javier Picorel, and Vasileios Trigonakis, <i>École Polytechnique Fédérale de Lausanne (EPFL)</i>	
<b>Greening the Video Transcoding Service with Low-Cost Hardware Transcoders</b> .....	<b>407</b>
Peng Liu, <i>University of Wisconsin—Madison</i> ; Jongwon Yoon, <i>Hanyang University</i> ; Lance Johnson, <i>University of Minnesota</i> ; Suman Banerjee, <i>University of Wisconsin—Madison</i>	
<b>MEANTIME: Achieving Both Minimal Energy and Timeliness with Approximate Computing</b> .....	<b>421</b>
Anne Farrell and Henry Hoffmann, <i>University of Chicago</i>	
<b>Network Design and Usage Studies</b>	
<b>Design Guidelines for High Performance RDMA Systems</b> .....	<b>437</b>
Anuj Kalia, <i>Carnegie Mellon University</i> ; Michael Kaminsky, <i>Intel Labs</i> ; David G. Andersen, <i>Carnegie Mellon University</i>	
<b>Balancing CPU and Network in the Cell Distributed B-Tree Store</b> .....	<b>451</b>
Christopher Mitchell, Kate Montgomery, and Lamont Nelson, <i>New York University</i> ; Siddhartha Sen, <i>Microsoft Research</i> ; Jinyang Li, <i>New York University</i>	
<b>An Evolutionary Study of Linux Memory Management for Fun and Profit</b> .....	<b>465</b>
Jian Huang, Moinuddin K. Qureshi, and Karsten Schwan, <i>Georgia Institute of Technology</i>	
<b>Getting Back Up: Understanding How Enterprise Data Backups Fail</b> .....	<b>479</b>
George Amvrosiadis, <i>University of Toronto</i> ; Medha Bhadkamkar, <i>Veritas Labs</i>	

# Friday, June 24, 2016

## Data Is Now Big Data

**SplitJoin: A Scalable, Low-latency Stream Join Architecture with Adjustable Ordering Precision** .....493  
Mohammadreza Najafi, *Technische Universität München*; Mohammad Sadoghi, *IBM T. J. Watson Research Center*; Hans-Arno Jacobsen, *Middleware Systems Research Group*

**Load the Edges You Need: A Generic I/O Optimization for Disk-based Graph Processing** .....507  
Keval Vora, *University of California, Riverside*; Guoqing Xu, *University of California, Irvine*; Rajiv Gupta, *University of California, Riverside*

**Version Traveler: Fast and Memory-Efficient Version Switching in Graph Processing Systems** .....523  
Xiaoen Ju, *University of Michigan*; Dan Williams and Hani Jamjoom, *IBM T. J. Watson Research Center*; Kang G. Shin, *University of Michigan*

**Tucana: Design and Implementation of a Fast and Efficient Scale-up Key-value Store** .....537  
Anastasios Papagiannis, *Foundation of Research and Technology-Hellas (FORTH) and University of Crete*; Giorgos Saloustros, *Foundation of Research and Technology-Hellas (FORTH)*; Pilar González-Férez, *Foundation of Research and Technology-Hellas (FORTH) and University of Murcia*; Angelos Bilas, *Foundation of Research and Technology-Hellas (FORTH) and University of Crete*

## Virtualization

**Samsara: Efficient Deterministic Replay in Multiprocessor Environments with Hardware Virtualization Extensions** .....551  
Shiru Ren, Le Tan, Chunqi Li, and Zhen Xiao, *Peking University*; Weijia Song, *Cornell University*

**Hardware-Assisted On-Demand Hypervisor Activation for Efficient Security Critical Code Execution on Mobile Devices** .....565  
Yeongpil Cho, *Seoul National University*; Junbum Shin, *Samsung Electronics*; Donghyun Kwon, *Seoul National University*; MyungJoo Ham and Yuna Kim, *Samsung Electronics*; Yunheung Paek, *Seoul National University*

**gScale: Scaling up GPU Virtualization with Dynamic Sharing of Graphics Memory Space** .....579  
Mochi Xue, *Shanghai Jiao Tong University and Intel Corporation*; Kun Tian, *Intel Corporation*; Yaozu Dong, *Shanghai Jiao Tong University and Intel Corporation*; Jiacheng Ma, Jiajun Wang, and Zhengwei Qi, *Shanghai Jiao Tong University*; Bingsheng He, *National University of Singapore*; Haibing Guan, *Shanghai Jiao Tong University*

**A General Persistent Code Caching Framework for Dynamic Binary Translation (DBT)** .....591  
Wenwen Wang, Pen-Chung Yew, Antonia Zhai, and Stephen McCamant, *University of Minnesota, Twin Cities*

## Operating Systems

**Instant OS Updates via Userspace Checkpoint-and-Restart** .....605  
Sanidhya Kashyap, Changwoo Min, Byoungyoung Lee, and Taesoo Kim, *Georgia Institute of Technology*; Pavel Emelyanov, *CRIU and Odin, Inc.*

**Apps with Hardware: Enabling Run-time Architectural Customization in Smart Phones** .....621  
Michael Coughlin, Ali Ismail, and Eric Keller, *University of Colorado, Boulder*

**Testing Error Handling Code in Device Drivers Using Characteristic Fault Injection** .....635  
Jia-Ju Bai, Yu-Ping Wang, Jie Yin, and Shi-Min Hu, *Tsinghua University*

**Multicore Locks: The Case Is Not Closed Yet** .....649  
Hugo Guiroux and Renaud Lachaize, *Université Grenoble Alpes and Laboratoire d'Informatique de Grenoble*; Vivien Quéma, *Université Grenoble Alpes, Grenoble Institute of Technology, and Laboratoire d'Informatique de Grenoble*