

# 2020 USENIX Annual Technical Conference

July 15–17, 2020

USENIX ATC '20 Program Co-Chairs Message..... x

## Wednesday, July 15

### The Non-Volatile One

**Libnvmio: Reconstructing Software IO Path with Failure-Atomic Memory-Mapped Interface** ..... 1

Jungsik Choi, *Sungkyunkwan University*; Jaewan Hong and Youngjin Kwon, *KAIST*; Hwansoo Han, *Sungkyunkwan University*

**MatrixKV: Reducing Write Stalls and Write Amplification in LSM-tree Based KV Stores with Matrix Container in NVM** .....17

Ting Yao, Yiwen Zhang, and Jiguang Wan, *Huazhong University of Science and Technology*; Qiu Cui and Liu Tang, *PingCAP*; Hong Jiang, *UT Arlington*; Changsheng Xie, *Huazhong University of Science and Technology*; Xubin He, *Temple University*

**Disaggregating Persistent Memory and Controlling Them Remotely: An Exploration of Passive Disaggregated Key-Value Stores** ..... 33

Shin-Yeh Tsai, *Purdue University*; Yizhou Shan and Yiying Zhang, *University of California, San Diego*

**SplinterDB: Closing the Bandwidth Gap for NVMe Key-Value Stores**..... 49

Alexander Conway, *Rutgers University and VMware Research*; Abhishek Gupta, *DropBox*; Vijay Chidambaram, *University of Texas at Austin and VMware Research*; Martin Farach-Colton, *Rutgers University*; Richard Spillane, *VMware*; Amy Tai and Rob Johnson, *VMware Research*

**Twizzler: a Data-Centric OS for Non-Volatile Memory** ..... 65

Daniel Bittman and Peter Alvaro, *UC Santa Cruz*; Pankaj Mehra, *IEEE Member*; Darrell D. E. Long, *UC Santa Cruz*; Ethan L. Miller, *UC Santa Cruz / Pure Storage*

### The Data Center One

**BASTION: A Security Enforcement Network Stack for Container Networks**..... 81

Jaehyun Nam, Seungsoo Lee, and Hyunmin Seo, *KAIST*; Phil Porras and Vinod Yegneswaran, *SRI International*; Seungwon Shin, *KAIST*

**Spool: Reliable Virtualized NVMe Storage Pool in Public Cloud Infrastructure** ..... 97

Shuai Xue, Shang Zhao, and Quan Chen, *Shanghai Jiao Tong University and Alibaba Cloud*; Gang Deng, Zheng Liu, Jie Zhang, Zhuo Song, Tao Ma, Yong Yang, Yanbo Zhou, Keqiang Niu, and Sijie Sun, *Alibaba Cloud*; Minyi Guo, *Shanghai Jiao Tong University*

**HDDse: Enabling High-Dimensional Disk State Embedding for Generic Failure Detection System of Heterogeneous Disks in Large Data Centers** ..... 111

Ji Zhang, *Huazhong University of Science and Technology and University of Amsterdam*; Ping Huang, *Huazhong University of Science and Technology and Temple University*; Ke Zhou, *Huazhong University of Science and Technology*; Ming Xie, *Tencent Inc.*; Sebastian Schelter, *University of Amsterdam*

**Adaptive Placement for In-memory Storage Functions** ..... 127

Ankit Bhardwaj, Chinmay Kulkarni, and Ryan Stutsman, *University of Utah*

**NetKernel: Making Network Stack Part of the Virtualized Infrastructure**..... 143

Zhixiong Niu, *Microsoft Research*; Hong Xu, *City University of Hong Kong*; Peng Cheng, *Microsoft Research*; Qiang Su, *City University of Hong Kong*; Yongqiang Xiong, *Microsoft Research*; Tao Wang, *New York University*; Dongsu Han, *KAIST*; Keith Winstein, *Stanford University*

## The Cloudy One

- Platinum: A CPU-Efficient Concurrent Garbage Collector for Tail-Reduction of Interactive Services** . . . . . 159  
Mingyu Wu, Ziming Zhao, Yanfei Yang, Haoyu Li, Haibo Chen, Binyu Zang, and Haibing Guan, *Shanghai Jiao Tong University*; Sanhong Li, Chuansheng Lu, and Tongbao Zhang, *Alibaba*
- PinK: High-speed In-storage Key-value Store with Bounded Tails** . . . . . 173  
Junsu Im and Jinwook Bae, *DGIST*; Chanwoo Chung and Arvind, *Massachusetts Institute of Technology*; Sungjin Lee, *DGIST*
- OPTIMUSCLOUD: Heterogeneous Configuration Optimization for Distributed Databases in the Cloud** . . . . . 189  
Ashraf Mahgoub and Alexander Michaelson Medoff, *Purdue University*; Rakesh Kumar, *Microsoft*; Subrata Mitra, *Adobe Research*; Ana Klimovic, *Google Research*; Somali Chaterji and Saurabh Bagchi, *Purdue University*
- Serverless in the Wild: Characterizing and Optimizing the Serverless Workload at a Large Cloud Provider** . . . . . 205  
Mohammad Shahrad, Rodrigo Fonseca, Íñigo Goiri, Gohar Chaudhry, Paul Batum, Jason Cooke, Eduardo Laureano, Colby Tresness, Mark Russinovich, and Ricardo Bianchini, *Microsoft Azure and Microsoft Research*
- Lessons Learned from the Chameleon Testbed** . . . . . 219  
Kate Keahey, *Argonne National Laboratory*; Jason Anderson and Zhuo Zhen, *University of Chicago*; Pierre Riteau, *StackHPC Ltd*; Paul Ruth, *RENCI UNC Chapel Hill*; Dan Stanzione, *Texas Advanced Computing Center*; Mert Cevik, *RENCI UNC Chapel Hill*; Jacob Colleran and Haryadi S. Gunawi, *University of Chicago*; Cody Hammock, *Texas Advanced Computing Center*; Joe Mambretti, *Northwestern University*; Alexander Barnes, François Halbah, Alex Rocha, and Joe Stubbs, *Texas Advanced Computing Center*

## The Buggy One

- SPINFER: Inferring Semantic Patches for the Linux Kernel** . . . . . 235  
Lucas Serrano and Van-Anh Nguyen, *Sorbonne University/Inria/LIP6*; Ferdian Thung, Lingxiao Jiang, and David Lo, *School of Information Systems, Singapore Management University*; Julia Lawall and Gilles Muller, *Inria/Sorbonne University/LIP6*
- FuZZan: Efficient Sanitizer Metadata Design for Fuzzing** . . . . . 249  
Yuseok Jeon, *Purdue University*; WookHyun Han, *KAIST*; Nathan Burow, *Purdue University*; Mathias Payer, *EPFL*
- PracExtractor: Extracting Configuration Good Practices from Manuals to Detect Server Misconfigurations** . . . . . 265  
Chengcheng Xiang and Haochen Huang, *University of California San Diego*; Andrew Yoo, *University of Illinois at Urbana-Champaign*; Yuanyuan Zhou, *University of California, San Diego*; Shankar Pasupathy, *NetApp*
- Reverse Debugging of Kernel Failures in Deployed Systems** . . . . . 281  
Xinyang Ge, *Microsoft Research*; Ben Niu, *Microsoft*; Weidong Cui, *Microsoft Research*
- Offload Annotations: Bringing Heterogeneous Computing to Existing Libraries and Workloads** . . . . . 293  
Gina Yuan, Shoumik Palkar, Deepak Narayanan, and Matei Zaharia, *Stanford University*

## Thursday, July 16

### The Machine Learning One

- HetPipe: Enabling Large DNN Training on (Whimpy) Heterogeneous GPU Clusters through Integration of Pipelined Model Parallelism and Data Parallelism** . . . . . 307  
Jay H. Park, Gyeongchan Yun, Chang M. Yi, Nguyen T. Nguyen, and Seungmin Lee, *UNIST*; Jaesik Choi, *KAIST*; Sam H. Noh and Young-ri Choi, *UNIST*
- AutoSys: The Design and Operation of Learning-Augmented Systems** . . . . . 323  
Chieh-Jan Mike Liang, Hui Xue, Mao Yang, and Lidong Zhou, *Microsoft Research*; Lifei Zhu, *Peking University and Microsoft Research*; Zhao Lucis Li and Zibo Wang, *University of Science and Technology of China and Microsoft Research*; Qi Chen and Quanlu Zhang, *Microsoft Research*; Chuanjie Liu, *Microsoft Bing Platform*; Wenjun Dai, *Microsoft Bing Ads*
- Daydream: Accurately Estimating the Efficacy of Performance Optimizations for DNN Training** . . . . . 337  
Hongyu Zhu, *University of Toronto & Vector Institute*; Amar Phanishayee, *Microsoft Research*; Gennady Pekhimenko, *University of Toronto & Vector Institute*

<b>ALERT: Accurate Learning for Energy and Timeliness</b> .....	<b>353</b>
Chengcheng Wan, Muhammad Santriaji, Eri Rogers, Henry Hoffmann, Michael Maire, and Shan Lu, <i>University of Chicago</i>	
<b>NeuOS: A Latency-Predictable Multi-Dimensional Optimization Framework for DNN-driven Autonomous Systems</b> .....	<b>371</b>
Soroush Bateni and Cong Liu, <i>University of Texas at Dallas</i>	
<b>PERCIVAL: Making In-Browser Perceptual Ad Blocking Practical with Deep Learning</b> .....	<b>387</b>
Zainul Abi Din, <i>UC Davis</i> ; Panagiotis Tigas, <i>University of Oxford</i> ; Samuel T. King, <i>UC Davis, Bouncer Technologies</i> ; Benjamin Livshits, <i>Brave Software, Imperial College London</i>	
<b>The OS and Virtualization One</b>	
<b>Harmonizing Performance and Isolation in Microkernels with Efficient Intra-kernel Isolation and Communication</b>	<b>401</b>
Jinyu Gu, Xinyue Wu, Wentai Li, Nian Liu, Zeyu Mi, Yubin Xia, and Haibo Chen, <i>Shanghai Jiao Tong University</i>	
<b>FAASM: Lightweight Isolation for Efficient Stateful Serverless Computing</b> .....	<b>419</b>
Simon Shillaker and Peter Pietzuch, <i>Imperial College London</i>	
<b>Fewer Cores, More Hertz: Leveraging High-Frequency Cores in the OS Scheduler for Improved Application Performance</b> .....	<b>435</b>
Redha Gouicem and Damien Carver, <i>Sorbonne University, LIP6, Inria</i> ; Jean-Pierre Lozi, <i>Oracle Labs</i> ; Julien Sopena, <i>Sorbonne University, LIP6, Inria</i> ; Baptiste Lepers and Willy Zwaenepoel, <i>University of Sydney</i> ; Nicolas Palix, <i>Université Grenoble Alpes</i> ; Julia Lawall and Gilles Muller, <i>Inria, Sorbonne University, LIP6</i>	
<b>vSMT-IO: Improving I/O Performance and Efficiency on SMT Processors in Virtualized Clouds</b> .....	<b>449</b>
Weiwei Jia, <i>New Jersey Institute of Technology</i> ; Jianchen Shan, <i>Hofstra University</i> ; Tsz On Li, <i>University of Hong Kong</i> ; Xiaowei Shang, <i>New Jersey Institute of Technology</i> ; Heming Cui, <i>University of Hong Kong</i> ; Xiaoning Ding, <i>New Jersey Institute of Technology</i>	
<b>Lightweight Preemptible Functions</b> .....	<b>465</b>
Sol Boucher, <i>Carnegie Mellon University</i> ; Anuj Kalia, <i>Microsoft Research</i> ; David G. Andersen, <i>Carnegie Mellon University</i> ; Michael Kaminsky, <i>BrdgAI / Carnegie Mellon University</i>	
<b>coIOMMU: A Virtual IOMMU with Cooperative DMA Buffer Tracking for Efficient Memory Management in Direct I/O</b> .....	<b>479</b>
Kun Tian, Yu Zhang, Luwei Kang, Yan Zhao, and Yaozu Dong, <i>Intel Corporation</i>	
<b>The WAN One</b>	
<b>BatchCrypt: Efficient Homomorphic Encryption for Cross-Silo Federated Learning</b> .....	<b>493</b>
Chengliang Zhang, Suyi Li, Junzhe Xia, and Wei Wang, <i>Hong Kong University of Science and Technology</i> ; Feng Yan, <i>University of Nevada, Reno</i> ; Yang Liu, <i>WeBank</i>	
<b>A Deep Dive into DNS Query Failures</b> .....	<b>507</b>
Donghui Yang, <i>Institute of Computing Technology, Chinese Academy of Sciences</i> ; Zhenyu Li, <i>Institute of Computing Technology, Chinese Academy of Sciences, and Purple Mountain Laboratories</i> ; Gareth Tyson, <i>Queen Mary University of London</i>	
<b>A Decentralized Blockchain with High Throughput and Fast Confirmation</b> .....	<b>515</b>
Chenxin Li, Peilun Li, and Dong Zhou, <i>Tsinghua University</i> ; Zhe Yang, Ming Wu, and Guang Yang, <i>Conflux Foundation</i> ; Wei Xu, <i>Tsinghua University</i> ; Fan Long, <i>University of Toronto and Conflux Foundation</i> ; Andrew Chi-Chih Yao, <i>Tsinghua University</i>	
<b>Reconstructing proprietary video streaming algorithms</b> .....	<b>529</b>
Maximilian Grüner, Melissa Licciardello, and Ankit Singla, <i>ETH Zürich</i>	
<b>Midgress-aware traffic provisioning for content delivery</b> .....	<b>543</b>
Aditya Sundarrajan, <i>University of Massachusetts Amherst</i> ; Mangesh Kasbekar, <i>Akamai Technologies</i> ; Ramesh K. Sitaraman, <i>University of Massachusetts Amherst &amp; Akamai Technologies</i> ; Samta Shukla, <i>CVS Health</i>	

## The One about Big Data

- GraphWalker: An I/O-Efficient and Resource-Friendly Graph Analytic System for Fast and Scalable Random Walks** ..... 559  
Rui Wang and Yongkun Li, *University of Science and Technology of China*; Hong Xie, *Chongqing University*; Yinlong Xu, *University of Science and Technology of China*; John C. S. Lui, *The Chinese University of Hong Kong*
- Scaph: Scalable GPU-Accelerated Graph Processing with Value-Driven Differential Scheduling** ..... 573  
Long Zheng, Xianliang Li, Yaohui Zheng, Yu Huang, Xiaofei Liao, and Hai Jin, *Huazhong University of Science and Technology*; Jingling Xue, *UNSW Sydney*; Zhiyuan Shao and Qiang-Sheng Hua, *Huazhong University of Science and Technology*
- Peregreen – modular database for efficient storage of historical time series in cloud environments** ..... 589  
Alexander Visheratin, Alexey Struckov, Semen Yufa, Alexey Muratov, Denis Nasonov, and Nikolay Butakov, *ITMO University*; Yury Kuznetsov and Michael May, *Siemens*
- AC-Key: Adaptive Caching for LSM-based Key-Value Stores** ..... 603  
Fenggang Wu, Ming-Hong Yang, Baoquan Zhang, and David H.C. Du, *University of Minnesota*
- POSH: A Data-Aware Shell** .....617  
Deepti Raghavan, Sadjad Fouladi, Philip Levis, and Matei Zaharia, *Stanford University*

## Friday, July 17

### The One about Acceleration

- FineStream: Fine-Grained Window-Based Stream Processing on CPU-GPU Integrated Architectures** ..... 633  
Feng Zhang and Lin Yang, *Renmin University of China*; Shuhao Zhang, *Technische Universität Berlin and National University of Singapore*; Bingsheng He, *National University of Singapore*; Wei Lu and Xiaoyong Du, *Renmin University of China*
- OpenExpress: Fully Hardware Automated Open Research Framework for Future Fast NVMe Devices** ..... 649  
Myoungsoo Jung, *KAIST*
- Fast Software Cache Design for Network Appliances** ..... 657  
Dong Zhou, *Tsinghua University*; Huacheng Yu, *Princeton University*; Michael Kaminsky, *BrdgAI*; David Andersen, *BrdgAI and Carnegie Mellon University*
- Reexamining Direct Cache Access to Optimize I/O Intensive Applications for Multi-hundred-gigabit Networks** .. 673  
Alireza Farshin, *KTH Royal Institute of Technology*; Amir Roozbeh, *KTH Royal Institute of Technology and Ericsson Research*; Gerald Q. Maguire Jr. and Dejan Kostić, *KTH Royal Institute of Technology*
- sRDMA – Efficient NIC-based Authentication and Encryption for Remote Direct Memory Access** ..... 691  
Konstantin Taranov, Benjamin Rothenberger, Adrian Perrig, and Torsten Hoefler, *ETH Zurich*
- UREQA: Leveraging Operation-Aware Error Rates for Effective Quantum Circuit Mapping on NISQ-Era Quantum Computers** ..... 705  
Tirthak Patel, Baolin Li, Rohan Basu Roy, and Devesh Tiwari, *Northeastern University*

### The One about Storage

- Austere Flash Caching with Deduplication and Compression** ..... 713  
Qiuping Wang and Jinhong Li, *The Chinese University of Hong Kong*; Wen Xia, *Harbin Institute of Technology, Shenzhen*; Erik Kruus and Biplob Debnath, *NEC Labs*; Patrick P. C. Lee, *The Chinese University of Hong Kong*
- DADI: Block-Level Image Service for Agile and Elastic Application Deployment** ..... 727  
Huiba Li, Yifan Yuan, Rui Du, Kai Ma, Lanzheng Liu, and Windsor Hsu, *Alibaba Group*
- Efficient Miss Ratio Curve Computation for Heterogeneous Content Popularity** .....741  
Damiano Carra, *University of Verona, Italy*; Giovanni Neglia, *Inria, Université Côte d'Azur, France*
- Can Applications Recover from fsync Failures?** ..... 753  
Anthony Rebello, Yuvraj Patel, Ramnathan Alagappan, Andrea C. Arpaci-Dusseau, and Remzi H. Arpaci-Dusseau, *University of Wisconsin - Madison*

**DupHunter: Flexible High-Performance Deduplication for Docker Registries** . . . . . 769  
Nannan Zhao, Hadeel Albahar, Subil Abraham, and Keren Chen, *Virginia Tech*; Vasily Tarasov, Dimitrios Skourtis, Lukas Rupperecht, and Ali Anwar, *IBM Research—Almaden*; Ali R. Butt, *Virginia Tech*

**OSCA: An Online-Model Based Cache Allocation Scheme in Cloud Block Storage Systems** . . . . . 785  
Yu Zhang, *Huazhong University of Science and Technology*; Ping Huang, *Huazhong University of Science and Technology and Temple University*; Ke Zhou and Hua Wang, *Huazhong University of Science and Technology*; Jianying Hu, Yongguang Ji, and Bin Cheng, *Tencent Inc.*

## **The Memorable One**

**Lock-free Concurrent Level Hashing for Persistent Memory** . . . . . 799  
Zhangyu Chen, Yu Hua, Bo Ding, and Pengfei Zuo, *Huazhong University of Science and Technology*

**Optimizing Memory-mapped I/O for Fast Storage Devices** . . . . . 813  
Anastasios Papagiannis, Giorgos Xanthakis, Giorgos Saloustros, Manolis Marazakis, and Angelos Bilas, *FORTH-ICS*

**A Comprehensive Analysis of Superpage Management Mechanisms and Policies** . . . . . 829  
Weixi Zhu, Alan L. Cox, and Scott Rixner, *Rice University*

**Effectively Prefetching Remote Memory with Leap** . . . . . 843  
Hasan Al Maruf and Mosharaf Chowdhury, *University of Michigan*

**go-pmem: Native Support for Programming Persistent Memory in Go** . . . . . 859  
Jerrin Shaji George, Mohit Verma, Rajesh Venkatasubramanian, and Pratap Subrahmanyam, *VMware*

**End the Senseless Killing: Improving Memory Management for Mobile Operating Systems** . . . . . 873  
Niel Lebeck, Arvind Krishnamurthy, and Henry M. Levy, *University of Washington*; Irene Zhang, *Microsoft Research*

## **The One on the Edge**

**Retnetwork: Exploring Reader Network with COTS RFID Systems** . . . . . 889  
Jia Liu and Xingyu Chen, *Nanjing University*; Shigang Chen, *University of Florida*; Wei Wang, Dong Jiang, and Lijun Chen, *Nanjing University*

**Acclaim: Adaptive Memory Reclaim to Improve User Experience in Android Systems** . . . . . 897  
Yu Liang and Jinheng Li, *City University of Hong Kong*; Rachata Ausavarungnirun, *King Mongkut's University of Technology North Bangkok*; Riwei Pan, *City University of Hong Kong*; Liang Shi, *East China Normal University*; Tei-Wei Kuo, *City University of Hong Kong and National Taiwan University*; Chun Jason Xue, *City University of Hong Kong*

**SweynTooth: Unleashing Mayhem over Bluetooth Low Energy** . . . . . 911  
Matheus E. Garbelini, *Singapore University of Technology and Design*; Chundong Wang, *ShanghaiTech University*; Sudipta Chattopadhyay, *Singapore University of Technology and Design*; Sun Sumei and Ernest Kurniawan, *A\*Star*

**Fine-Grained Isolation for Scalable, Dynamic, Multi-tenant Edge Clouds** . . . . . 927  
Yuxin Ren, *The George Washington University*; Guyue Liu, *Carnegie Mellon University*; Vlad Nitu, *INSA Lyon France*; Wenyan Shao, Riley Kennedy, Gabriel Parmer, and Timothy Wood, *The George Washington University*; Alain Tchana, *ENS Lyon France*

**Firefly: Untethered Multi-user VR for Commodity Mobile Devices** . . . . . 943  
Xing Liu, *University of Minnesota, Twin Cities*; Christina Vlachou, *Hewlett Packard Labs*; Feng Qian and Chendong Wang, *University of Minnesota, Twin Cities*; Kyu-Han Kim, *Hewlett Packard Labs*