

2022 USENIX Annual Technical Conference

July 11–13, 2022

Monday, July 11

Storage 1

ZNSwap: un-Block your Swap 1
Shai Bergman, *Technion*; Niklas Cassel and Matias Bjørling, *Western Digital*; Mark Silberstein, *Technion*

Building a High-performance Fine-grained Deduplication Framework for Backup Storage with High Deduplication Ratio 19
Xiangyu Zou and Wen Xia, *Harbin Institute of Technology, Shenzhen*; Philip Shilane, *Dell Technologies*; Haijun Zhang and Xuan Wang, *Harbin Institute of Technology, Shenzhen*

Secure and Lightweight Deduplicated Storage via Shielded Deduplication-Before-Encryption 37
Zuoru Yang, *The Chinese University of Hong Kong*; Jingwei Li, *University of Electronic Science and Technology of China*; Patrick P. C. Lee, *The Chinese University of Hong Kong*

Containers

RunD: A Lightweight Secure Container Runtime for High-density Deployment and High-concurrency Startup in Serverless Computing 53
Zijun Li, *Department of Computer Science and Engineering, Shanghai Jiao Tong University and Alibaba Group*; Jiagan Cheng, and Quan Chen, *Department of Computer Science and Engineering, Shanghai Jiao Tong University*; Eryu Guan, Zizheng Bian, Yi Tao, Bin Zha, Qiang Wang, and Weidong Han, *Alibaba Group*; Minyi Guo, *Department of Computer Science and Engineering, Shanghai Jiao Tong University*

Help Rather Than Recycle: Alleviating Cold Startup in Serverless Computing Through Inter-Function Container Sharing 69
Zijun Li, Linsong Guo, Quan Chen, Jiagan Cheng, and Chuhao Xu, *Shanghai Jiao Tong University*; Deze Zeng, *School of Computer Science, China University of Geosciences*; Zhuo Song, Tao Ma, and Yong Yang, *Alibaba Cloud*; Chao Li and Minyi Guo, *Department of Computer Science and Engineering, Shanghai Jiao Tong University*

RRC: Responsive Replicated Containers 85
Diyu Zhou, *UCLA and EPFL*; Yuval Tamir, *UCLA*

Distributed Systems 1

uKharon: A Membership Service for Microsecond Applications 101
Rachid Guerraoui and Antoine Murat, *École Polytechnique Fédérale de Lausanne (EPFL)*; Javier Picorel, *Huawei Technologies*; Athanasios Xygkis, *EPFL*; Huabing Yan and Pengfei Zuo, *Huawei Technologies*

KRCORE: A Microsecond-scale RDMA Control Plane for Elastic Computing 121
Xingda Wei, *Institute of Parallel and Distributed Systems, SEIEE, Shanghai Jiao Tong University, and Shanghai AI Laboratory*; Fangming Lu, *Institute of Parallel and Distributed Systems, SEIEE, Shanghai Jiao Tong University*; Rong Chen, *Institute of Parallel and Distributed Systems, SEIEE, Shanghai Jiao Tong University, and Shanghai AI Laboratory*; Haibo Chen, *Institute of Parallel and Distributed Systems, SEIEE, Shanghai Jiao Tong University*

Zero-Change Object Transmission for Distributed Big Data Analytics 137
Mingyu Wu, *Institute of Parallel and Distributed Systems, SEIEE, Shanghai Jiao Tong University and Shanghai AI Laboratory*; Shuaiwei Wang, *Institute of Parallel and Distributed Systems, SEIEE, Shanghai Jiao Tong University*; Haibo Chen and Binyu Zang, *Institute of Parallel and Distributed Systems, SEIEE, Shanghai Jiao Tong University and Engineering Research Center for Domain-specific Operating Systems, Ministry of Education, China*

Sift: Using Refinement-guided Automation to Verify Complex Distributed Systems 151
Haojun Ma, Hammad Ahmad, Aman Goel, Eli Goldweber, Jean-Baptiste Jeannin, Manos Kapritsos, and Baris Kasikci, *University of Michigan*

Machine Learning 1

Faith: An Efficient Framework for Transformer Verification on GPUs 167
Boyuan Feng, Tianqi Tang, Yuke Wang, Zhaodong Chen, Zheng Wang, Shu Yang, Yuan Xie, and Yufei Ding, *University of California, Santa Barbara*

DVABatch: Diversity-aware Multi-Entry Multi-Exit Batching for Efficient Processing of DNN Services on GPUs . 183
Weihao Cui, Han Zhao, Quan Chen, Hao Wei, and Zirui Li, *Shanghai Jiao Tong University*; Deze Zeng, *China University of Geosciences*; Chao Li and Minyi Guo, *Shanghai Jiao Tong University*

Serving Heterogeneous Machine Learning Models on Multi-GPU Servers with Spatio-Temporal Sharing 199
Seungbeom Choi, Sunho Lee, Yeonjae Kim, Jongse Park, Youngjin Kwon, and Jaehyuk Huh, *KAIST*

PilotFish: Harvesting Free Cycles of Cloud Gaming with Deep Learning Training 217
Wei Zhang and Binghao Chen, *Shanghai Jiao Tong University*; Zhenhua Han, *Microsoft Research Asia*; Quan Chen, *Shanghai Jiao Tong University*; Peng Cheng, Fan Yang, Ran Shu, and Yuqing Yang, *Microsoft Research*; Minyi Guo, *Shanghai Jiao Tong University*

Operating Systems 1

Privbox: Faster System Calls Through Sandboxed Privileged Execution 233
Dmitry Kuznetsov and Adam Morrison, *Tel Aviv University*

BBQ: A Block-based Bounded Queue for Exchanging Data and Profiling 249
Jiawei Wang, *Huawei Dresden Research Center, Huawei OS Kernel Lab and Technische Universität Dresden*; Diogo Behrens, Ming Fu, Lilith Oberhauser, Jonas Oberhauser, and Jitang Lei, *Huawei Dresden Research Center, Huawei OS Kernel Lab*; Geng Chen, *Huawei OS Kernel Lab*; Hermann Härtig, *Technische Universität Dresden*; Haibo Chen, *Huawei OS Kernel Lab and Shanghai Jiao Tong University*

Disaggregated Systems

Sibylla: To Retry or Not To Retry on Deep Learning Job Failure 263
Taeyoon Kim, Suyeon Jeong, Jongseop Lee, Soobee Lee, and Myeongjae Jeon, *UNIST*

Speculative Recovery: Cheap, Highly Available Fault Tolerance with Disaggregated Storage 271
Nanqinqin Li, Anja Kalaba, Michael J. Freedman, Wyatt Lloyd, and Amit Levy, *Princeton University*

Direct Access, High-Performance Memory Disaggregation with DirectCXL 287
Donghyun Gouk, Sangwon Lee, Miryeong Kwon, and Myoungsoo Jung, *Computer Architecture and Memory Systems Laboratory, Korea Advanced Institute of Science and Technology (KAIST)*

Networking 1

Not that Simple: Email Delivery in the 21st Century 295
Florian Holzbauer, *SBA Research*; Johanna Ullrich, *University of Vienna*; Martina Lindorfer, *TU Wien*; Tobias Fiebig, *Max-Planck-Institut für Informatik*

AddrMiner: A Comprehensive Global Active IPv6 Address Discovery System 309
Guanglei Song, Jiahai Yang, Lin He, Zhiliang Wang, *Institute for Network Sciences and Cyberspace, BNRist, Tsinghua University and Quan Cheng Laboratory, Jinan, Shandong, China*; Guo Li and Chenxin Duan, *Institute for Network Sciences and Cyberspace, BNRist, Tsinghua University*; Yaozhong Liu, *Tsinghua University*; Zhongxiang Sun, *School of Computer and Information Technology, Beijing Jiaotong University*

Co-opting Linux Processes for High-Performance Network Simulation 327
Rob Jansen, *U.S. Naval Research Laboratory*; Jim Newsome, *Tor Project*; Ryan Wails, *Georgetown University, U.S. Naval Research Laboratory*

Finding Bugs

KSG: Augmenting Kernel Fuzzing with System Call Specification Generation 351
Hao Sun, Yuheng Shen, Jianzhong Liu, Yiru Xu, and Yu Jiang, *Tsinghua University*

DLOS: Effective Static Detection of Deadlocks in OS Kernels 367
Jia-Ju Bai, Tuo Li, and Shi-Min Hu, *Tsinghua University*

Modulo: Finding Convergence Failure Bugs in Distributed Systems with Divergence Resync Models 383
Beom Heyn Kim, *Samsung Research and University of Toronto*; Taesoo Kim, *Samsung Research and Georgia Institute of Technology*; David Lie, *University of Toronto*

Tuesday, July 12

Security

SoftTRR: Protect Page Tables Against Rowhammer Attacks Using Software-Only Target Row Refresh **399**

Zhi Zhang, *CSIRO's Data61, Australia*; Yueqiang Cheng, *NIO Security Research*; Minghua Wang, *Baidu Security*; Wei He and Wenhao Wang, *State Key Laboratory of Information Security, Institute of Information Engineering, CAS, and University of Chinese Academy of Sciences*; Surya Nepal, *CSIRO's Data61, Australia*; Yansong Gao, *Nanjing University of Science and Technology, China*; Kang Li, *Baidu Security*; Zhe Wang and Chenggang Wu, *State Key Laboratory of Computer Architecture, Institute of Computing Technology, CAS, and University of Chinese Academy of Sciences*

Hardening Hypervisors with Ombro **415**

Ethan Johnson, Colin Pronovost, and John Criswell, *Department of Computer Science, University of Rochester*

HyperEnclave: An Open and Cross-platform Trusted Execution Environment **437**

Yuekai Jia, *Tsinghua University*; Shuang Liu, *Ant Group*; Wenhao Wang, *SKLOIS, Institute of Information Engineering, CAS, and School of Cyber Security, University of Chinese Academy of Sciences*; Yu Chen, *Tsinghua University*; Zhengde Zhai, Shoumeng Yan, and Zhengyu He, *Ant Group*

PRIDWEN: Universally Hardening SGX Programs via Load-Time Synthesis **455**

Fan Sang, *Georgia Institute of Technology*; Ming-Wei Shih, *Microsoft*; Sangho Lee, *Microsoft Research*; Xiaokuan Zhang, *Georgia Institute of Technology*; Michael Steiner and Mona Vij, *Intel Labs*; Taesoo Kim, *Georgia Institute of Technology*

Machine Learning 2

Tetris: Memory-efficient Serverless Inference through Tensor Sharing **473**

Jie Li, Laiping Zhao, and Yanan Yang, *College of Intelligence & Computing (CIC), Tianjin University, and Tianjin Key Lab of Advanced Networking (TANKLAB)*; Kunlin Zhan, *58.com*; Keqiu Li, *College of Intelligence & Computing (CIC), Tianjin University, and Tianjin Key Lab of Advanced Networking (TANKLAB)*

PetS: A Unified Framework for Parameter-Efficient Transformers Serving **489**

Zhe Zhou, *Peking University*; Xuechao Wei, *Peking University and Alibaba Group*; Jiejing Zhang, *Alibaba Group*; Guangyu Sun, *Peking University*

Campo: Cost-Aware Performance Optimization for Mixed-Precision Neural Network Training **505**

Xin He, *CSEE, Hunan University & Xidian University*; Jianhua Sun and Hao Chen, *CSEE, Hunan University*; Dong Li, *University of California, Merced*

Primo: Practical Learning-Augmented Systems with Interpretable Models **519**

Qinghao Hu, *Nanyang Technological University and S-Lab, NTU*; Harsha Nori, *Microsoft*; Peng Sun, *SenseTime*; Yonggang Wen and Tianwei Zhang, *Nanyang Technological University*

Distributed Systems 2

Meces: Latency-efficient Rescaling via Prioritized State Migration for Stateful Distributed Stream Processing Systems **539**

Rong Gu, Han Yin, Weichang Zhong, Chunfeng Yuan, and Yihua Huang, *State Key Laboratory for Novel Software Technology, Nanjing University*

DepFast: Orchestrating Code of Quorum Systems **557**

Xuhao Luo, *University of Illinois at Urbana-Champaign*; Weihai Shen and Shuai Mu, *Stony Brook University*; Tianyin Xu, *University of Illinois at Urbana-Champaign*

High Throughput Replication with Integrated Membership Management **575**

Pedro Fouto, Nuno Preguiça, and João Leitão, *NOVA LINCS & NOVA University Lisbon*

Operating Systems 2

CBMM: Financial Advice for Kernel Memory Managers **593**

Mark Mansi, Bijan Tabatabai, and Michael M. Swift, *University of Wisconsin - Madison*

EPK: Scalable and Efficient Memory Protection Keys **609**

Jinyu Gu, Hao Li, Wentai Li, Yubin Xia, and Haibo Chen, *Engineering Research Center for Domain-specific Operating Systems, Ministry of Education, China, Institute of Parallel and Distributed Systems (IPADS), SEIEE, Shanghai Jiao Tong University*

Memory Harvesting in Multi-GPU Systems with Hierarchical Unified Virtual Memory **625**

Sangjin Choi and Taeksoo Kim, *KAIST*; Jinwoo Jeong, *Ajou University*; Rachata Ausavarungnirun, *KMUTNB*; Myeongjae Jeon, *UNIST*; Youngjin Kwon, *KAIST*; Jeongseob Ahn, *Ajou University*

Deployed Systems 1

Zero Overhead Monitoring for Cloud-native Infrastructure using RDMA **639**

Zhe Wang, *Shanghai Jiao Tong University*; Teng Ma, *Alibaba Group*; Linghe Kong, *Shanghai Jiao Tong University*; Zhenzao Wen, Jingxuan Li, Zhuo Song, Yang Lu, Yong Yang, and Tao Ma, *Alibaba Group*; Guihai Chen, *Shanghai Jiao Tong University*; Wei Cao, *Alibaba Group*

CRISP: Critical Path Analysis of Large-Scale Microservice Architectures **655**

Zhizhou Zhang, *UC Santa Barbara*; Murali Krishna Ramanathan, Prithvi Raj, and Abhishek Parwal, *Uber Technologies Inc.*; Timothy Sherwood, *UC Santa Barbara*; Milind Chabbi, *Uber Technologies Inc.*

Whale: Efficient Giant Model Training over Heterogeneous GPUs **673**

Xianyan Jia, Le Jiang, Ang Wang, and Wencong Xiao, *Alibaba Group*; Ziji Shi, *Alibaba Group and National University of Singapore*; Jie Zhang, Xinyuan Li, Langshi Chen, Yong Li, Zhen Zheng, Xiaoyong Liu, and Wei Lin, *Alibaba Group*

Machine Learning 3

Cachew: Machine Learning Input Data Processing as a Service **689**

Dan Graur, Damien Aymon, Dan Kluser, and Tanguy Albrici, *ETH Zurich*; Chandramohan A. Thekkath, *Google*; Ana Klimovic, *ETH Zurich*

CoVA: Exploiting Compressed-Domain Analysis to Accelerate Video Analytics **707**

Jinwoo Hwang, Minsu Kim, Daeun Kim, Seungho Nam, Yoosung Kim, and Dohee Kim, *KAIST*; Hardik Sharma, *Google*; Jongse Park, *KAIST*

SOTER: Guarding Black-box Inference for General Neural Networks at the Edge **723**

Tianxiang Shen, Ji Qi, Jianyu Jiang, Xian Wang, Siyuan Wen, Xusheng Chen, and Shixiong Zhao, *The University of Hong Kong*; Sen Wang and Li Chen, *Huawei Technologies*; Xiapu Luo, *The Hong Kong Polytechnic University*; Fengwei Zhang, *Southern University of Science and Technology (SUSTech)*; Heming Cui, *The University of Hong Kong*

Storage 2

IPLFS: Log-Structured File System without Garbage Collection **739**

Juwon Kim, Minsu Jang, Muhammad Danish Tehseen, Joontaek Oh, and YouJip Won, *KAIST*

Vigil-KV: Hardware-Software Co-Design to Integrate Strong Latency Determinism into Log-Structured Merge Key-Value Stores **755**

Miryong Kwon, Seungjun Lee, and Hyunkyu Choi, *KAIST*; Jooyoung Hwang, *Samsung Electronics Co., Ltd.*; Myoungsoo Jung, *KAIST*

Pacman: An Efficient Compaction Approach for Log-Structured Key-Value Store on Persistent Memory **773**

Jing Wang, Youyou Lu, Qing Wang, and Minhui Xie, *Tsinghua University*; Keji Huang, *Huawei Technologies Co., Ltd*; Jiwu Shu, *Tsinghua University*

Networking 2

Towards Latency Awareness for Content Delivery Network Caching **789**

Gang Yan and Jian Li, *SUNY-Binghamton University*

Hashing Design in Modern Networks: Challenges and Mitigation Techniques **805**

Yunhong Xu, *Texas A&M University*; Keqiang He and Rui Wang, *Google*; Minlan Yu, *Harvard University and Google*; Nick Duffield, *Texas A&M University*; Hassan Wassel, Shidong Zhang, Leon Poutievski, Junlan Zhou, and Amin Vahdat, *Google*

Firebolt: Finding Bugs in Programmable Data Plane Generators **819**

Jiamin Cao, *Tsinghua University*; Yu Zhou and Chen Sun, *Alibaba Group*; Lin He, Zhaowei Xi, and Ying Liu, *Tsinghua University*

Wednesday, July 13

Compilers and PL

- Investigating Managed Language Runtime Performance: Why JavaScript and Python are 8x and 29x slower than C++, yet Java and Go can be Faster?** 835
David Lion, *University of Toronto and YScope Inc.*; Adrian Chiu and Michael Stumm, *University of Toronto*; Ding Yuan, *University of Toronto and YScope Inc.*
- Automatic Recovery of Fine-grained Compiler Artifacts at the Binary Level** 853
Yufei Du, *University of North Carolina at Chapel Hill*; Ryan Court and Kevin Snow, *Zeropoint Dynamics*; Fabian Monroe, *University of North Carolina at Chapel Hill*
- JITServer: Disaggregated Caching JIT Compiler for the JVM in the Cloud** 869
Alexey Khrabrov, *University of Toronto*; Marius Pirvu and Vijay Sundaesan, *IBM*; Eyal de Lara, *University of Toronto*
- Riker: Always-Correct and Fast Incremental Builds from Simple Specifications.** 885
Charlie Curtsinger, *Grinnell College*; Daniel W. Barowy, *Williams College*

Storage 3

- FlatFS: Flatten Hierarchical File System Namespace on Non-volatile Memories** 899
Miao Cai, *Key Laboratory of Water Big Data Technology of Ministry of Water Resources, Hohai University; School of Computer and Information, Hohai University; State Key Laboratory for Novel Software Technology, Nanjing University*; Junru Shen, *School of Computer and Information, Hohai University*; Bin Tang, *School of Computer and Information, Hohai University*; Hao Huang, *State Key Laboratory for Novel Software Technology, Nanjing University*; Baoliu Ye, *State Key Laboratory for Novel Software Technology, Nanjing University*; Key Laboratory of Water Big Data Technology of Ministry of Water Resources, Hohai University; School of Computer and Information, Hohai University
- StRAID: Stripe-threaded Architecture for Parity-based RAIDs with Ultra-fast SSDs** 915
Shucheng Wang, Qiang Cao, and Ziyi Lu, *Wuhan National Laboratory for Optoelectronics, HUST*; Hong Jiang, *Department of Computer Science and Engineering, UT Arlington*; Jie Yao, *School of Computer Science and Technology, HUST*; Yuanyuan Dong, *Alibaba Group*
- VINTER: Automatic Non-Volatile Memory Crash Consistency Testing for Full Systems** 933
Samuel Kalbfleisch, Lukas Werling, and Frank Bellosa, *Karlsruhe Institute of Technology*

NICs

- AINiCo: SmartNIC-accelerated Contention-aware Request Scheduling for Transaction Processing.** 951
Junru Li, Youyou Lu, Qing Wang, Jiazhen Lin, Zhe Yang, and Jiwu Shu, *Beijing National Research Center for Information Science and Technology (BNRist)*
- FpgaNIC: An FPGA-based Versatile 100Gb SmartNIC for GPUs** 967
Zeke Wang, Hongjing Huang, Jie Zhang, *Collaborative Innovation Center of Artificial Intelligence, Zhejiang University, China*; Fei Wu, *Collaborative Innovation Center of Artificial Intelligence, Zhejiang University, China, and Shanghai Institute for Advanced Study of Zhejiang University, China*; Gustavo Alonso, *ETH Zurich*
- Faster Software Packet Processing on FPGA NICs with eBPF Program Warping.** 987
Marco Bonola, *CNIT/Axbryd*; Giacomo Belocchi, Angelo Tulumello, and Marco Spaziani Brunella, *Axbryd/University of Rome Tor Vergata*; Giuseppe Siracusano, *NEC Laboratories Europe*; Giuseppe Bianchi, *University of Rome Tor Vergata*; Roberto Bifulco, *NEC Laboratories Europe*

Deployed Systems 2

- NVMe SSD Failures in the Field: the Fail-Stop and the Fail-Slow** 1005
Ruiming Lu, *Shanghai Jiao Tong University*; Erci Xu, *PDL*; Yiming Zhang, *Xiamen University*; Zhaosheng Zhu, Mengtian Wang, and Zongpeng Zhu, *Alibaba Inc.*; Guangtao Xue, *Shanghai Jiao Tong University*; Minglu Li, *Shanghai Jiao Tong University & Zhejiang Normal University*; Jiasheng Wu, *Alibaba Inc.*
- CacheSack: Admission Optimization for Google Datacenter Flash Caches.** 1021
Tzu-Wei Yang, Seth Pollen, Mustafa Uysal, Arif Merchant, and Homer Wolfmeister, *Google*
- Amazon DynamoDB: A Scalable, Predictably Performant, and Fully Managed NoSQL Database Service** 1037
Mostafa Elhemali, Niall Gallagher, Nicholas Gordon, Joseph Idziorek, Richard Krog, Colin Lazier, Erben Mo, Akhilesh Mritunjai, Somu Perianayagam, Tim Rath, Swami Sivasubramanian, James Christopher Sorenson III, Sroaj Sothikul, Doug Terry, Akshat Vig, *Amazon Web Services*