

2021 USENIX Annual Technical Conference

July 14–16, 2021

Message from the USENIX ATC '21 Program Co-Chairs x

Wednesday, July 14

Peeking over the Fence: RDMA

Naos: Serialization-free RDMA networking in Java 1
Konstantin Taranov, *ETH Zurich*; Rodrigo Bruno, *INESC-ID / Técnico, ULisboa*; Gustavo Alonso and Torsten Hoefler, *ETH Zurich*

One-sided RDMA-Conscious Extendible Hashing for Disaggregated Memory 15
Pengfei Zuo, Jiazhao Sun, Liu Yang, and Shuangwu Zhang, *Huawei Cloud*; Yu Hua, *Huazhong University of Science and Technology*

Characterizing and Optimizing Remote Persistent Memory with RDMA and NVM 31
Xingda Wei, Xiating Xie, Rong Chen, Haibo Chen, and Binyu Zang, *Shanghai Jiao Tong University; Shanghai AI Laboratory; Engineering Research Center for Domain-specific Operating Systems*

MigrOS: Transparent Live-Migration Support for Containerised RDMA Applications 47
Maksym Planeta and Jan Bierbaum, *TU Dresden*; Leo Sahaya Daphne Antony, *AMOLF*; Torsten Hoefler, *ETH Zürich*; Hermann Härtig, *TU Dresden*

Dogs Never Get Tired: Power and Edge Computing

Prediction-Based Power Oversubscription in Cloud Platforms 65
Alok Gautam Kumbhare, Reza Azimi, Ioannis Manousakis, Anand Bonde, Felipe Frujeri, Nithish Mahalingam, Pulkit A. Misra, Seyyed Ahmad Javadi, Bianca Schroeder, Marcus Fontoura, and Ricardo Bianchini, *Microsoft Research and Microsoft Azure*

Proactive Energy-Aware Adaptive Video Streaming on Mobile Devices 81
Jiayi Meng, Qiang Xu, and Y. Charlie Hu, *Purdue University*

Video Analytics with Zero-streaming Cameras 99
Mengwei Xu, *Peking University/Beijing University of Posts and Telecommunications*; Tiantu Xu, *Purdue ECE*; Yunxin Liu, *Institute for AI Industry Research (AIR), Tsinghua University*; Felix Xiaozhu Lin, *University of Virginia*

ASAP: Fast Mobile Application Switch via Adaptive Prepaging 117
Sam Son, Seung Yul Lee, Yunho Jin, and Jonghyun Bae, *Seoul National University*; Jinkyu Jeong, *Sungkyunkwan University*; Tae Jun Ham and Jae W. Lee, *Seoul National University*; Hongil Yoon, *Google*

Barking up the Wrong Tree: Correctness and Debugging

PyLIVE: On-the-Fly Code Change for Python-based Online Services 131
Haochen Huang, Chengcheng Xiang, Li Zhong, and Yuanyuan Zhou, *University of California, San Diego*

RIFF: Reduced Instruction Footprint for Coverage-Guided Fuzzing 147
Mingzhe Wang, Jie Liang, Chijin Zhou, and Yu Jiang, *Tsinghua University*; Rui Wang, *Capital Normal University*; Chengnian Sun, *Waterloo University*; Jiaguang Sun, *Tsinghua University*

TCP-Fuzz: Detecting Memory and Semantic Bugs in TCP Stacks with Fuzzing 161
Yong-Hao Zou and Jia-Ju Bai, *Tsinghua University*; Jielong Zhou, Jianfeng Tan, and Chenggang Qin, *Ant Group*; Shi-Min Hu, *Tsinghua University*

MLEE: Effective Detection of Memory Leaks on Early-Exit Paths in OS Kernels 177
Wenwen Wang, *University of Georgia*

Argus: Debugging Performance Issues in Modern Desktop Applications with Annotated Causal Tracing 193
Lingmei Weng, *Columbia University*; Peng Huang, *Johns Hopkins University*; Jason Nieh and Junfeng Yang, *Columbia University*

Searching for Tracks: Graphs

aDFS: An Almost Depth-First-Search Distributed Graph-Querying System 209
Vasileios Trigonakis and Jean-Pierre Lozi, *Oracle Labs*; Tomáš Faltín, *Oracle Labs and Charles University*; Nicholas P. Roth, *KUNGFU.AI*; Iraklis Psaroudakis, Arnaud Delamare, Vlad Haprian, Călin Iorgulescu, Petr Koupy, Jinsoo Lee, Sungpack Hong, and Hassan Chafi, *Oracle Labs*

GLIST: Towards In-Storage Graph Learning 225
Cangyuan Li, Ying Wang, Cheng Liu, and Shengwen Liang, *SKLCA, Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China; University of Chinese Academy of Sciences, Beijing, China*; Huawei Li, *SKLCA, Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China; University of Chinese Academy of Sciences, Beijing, China*; Peng Cheng Laboratory, *Shenzhen, China*; Xiaowei Li, *SKLCA, Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China; University of Chinese Academy of Sciences, Beijing, China*

DART: A Scalable and Adaptive Edge Stream Processing Engine 239
Pinchao Liu, *Florida International University*; Dilma Da Silva, *Texas A&M University*; Liting Hu, *Virginia Tech*

CrystalPerf: Learning to Characterize the Performance of Dataflow Computation through Code Analysis 253
Huangshi Tian, *HKUST*; Minchen Yu, *HKUST and Huawei Technologies Ltd.*; Wei Wang, *HKUST*

Controlling Memory Footprint of Stateful Streaming Graph Processing 269
Pourya Vaziri and Keval Vora, *Simon Fraser University*

Please Don't Chain Me Up: Blockchain and Security

Avocado: A Secure In-Memory Distributed Storage System 285
Maurice Bailleu, Dimitra Giantsidi, and Vasilis Gavrielatos, *University of Edinburgh*; Do Le Quoc, *Huawei Research*; Vijay Nagarajan, *University of Edinburgh*; Pramod Bhatotia, *University of Edinburgh and TU Munich*

Accelerating Encrypted Deduplication via SGX 303
Yanjing Ren and Jingwei Li, *University of Electronic Science and Technology of China*; Zuoru Yang and Patrick P. C. Lee, *The Chinese University of Hong Kong*; Xiaosong Zhang, *University of Electronic Science and Technology of China*

ICARUS: Attacking low Earth orbit satellite networks 317
Giacomo Giuliani, Tommaso Ciussani, Adrian Perrig, and Ankit Singla, *ETH Zurich*

RainBlock: Faster Transaction Processing in Public Blockchains 333
Soujanya Ponnappalli, Aashaka Shah, and Souvik Banerjee, *University of Texas at Austin*; Dahlia Malkhi, *Diem Association and Novi Financial*; Amy Tai, *VMware Research*; Vijay Chidambaram, *University of Texas at Austin and VMware Research*; Michael Wei, *VMware Research*

An Off-The-Chain Execution Environment for Scalable Testing and Profiling of Smart Contracts 349
Yeonsoo Kim and Seongho Jeong, *Yonsei University*; Kamil Jezek, *The University of Sydney*; Bernd Burgstaller, *Yonsei University*; Bernhard Scholz, *The University of Sydney*

I'm Old But I Learned a New Trick: Machine Learning

Octo: INT8 Training with Loss-aware Compensation and Backward Quantization for Tiny On-device Learning . 365
Qihua Zhou and Song Guo, *Hong Kong Polytechnic University*; Zhihao Qu, *Hohai University*; Jingcai Guo, Zhenda Xu, Jiewei Zhang, Tao Guo, and Boyuan Luo, *Hong Kong Polytechnic University*; Jingren Zhou, *Alibaba Group*

Fine-tuning giant neural networks on commodity hardware with automatic pipeline model parallelism 381
Saar Eliad, Ido Hakimi, and Alon De Jagger, *Department of Computer Science, Technion - Israel Institute of Technology*; Mark Silberstein, *Department of Computer Science and Department of Electrical Engineering, Technion - Israel Institute of Technology*; Assaf Schuster, *Department of Computer Science, Technion - Israel Institute of Technology*

INFaaS: Automated Model-less Inference Serving 397
Francisco Romero, Qian Li, Neeraja J. Yadwadkar, and Christos Kozyrakis, *Stanford University*

Jump-Starting Multivariate Time Series Anomaly Detection for Online Service Systems 413
Minghua Ma, *Tsinghua University, BNRist*; Shenglin Zhang, *Nankai University*; Junjie Chen, *Tianjin University*; Jim Xu, *Georgia Tech*; Haozhe Li and Yongliang Lin, *Nankai University*; Xiaohui Nie, *Tsinghua University, BNRist*; Bo Zhou and Yong Wang, *CNCERT/CC*; Dan Pei, *Tsinghua University, BNRist*

Palleon: A Runtime System for Efficient Video Processing toward Dynamic Class Skew 427
Boyuan Feng, Yuke Wang, Gushu Li, Yuan Xie, and Yufei Ding, *University of California, Santa Barbara*

Thursday, July 15

Can I Come In? It's Raining!: Cloud Computing

- FAASNET: Scalable and Fast Provisioning of Custom Serverless Container Runtimes at Alibaba Cloud Function Compute** 443
Ao Wang, *George Mason University*; Shuai Chang, *Alibaba Group*; Huangshi Tian, *Hong Kong University of Science and Technology*; Hongqi Wang, Haoran Yang, Huiba Li, and Rui Du, *Alibaba Group*; Yue Cheng, *George Mason University*
- Experiences in Managing the Performance and Reliability of a Large-Scale Genomics Cloud Platform** 459
Michael Hao Tong, Robert L. Grossman, and Haryadi S. Gunawi, *University of Chicago*
- Scaling Large Production Clusters with Partitioned Synchronization** 473
Yihui Feng, *Alibaba Group*; Zhi Liu, Yunjian Zhao, Tatiana Jin, and Yidi Wu, *The Chinese University of Hong Kong*; Yang Zhang, *Alibaba Group*; James Cheng, *The Chinese University of Hong Kong*; Chao Li and Tao Guan, *Alibaba Group*
- Fighting the Fog of War: Automated Incident Detection for Cloud Systems** 489
Liquan Li and Xu Zhang, *Microsoft Research*; Xin Zhao, *University of Chinese Academy of Sciences*; Hongyu Zhang, *The University of Newcastle*; Yu Kang, Pu Zhao, Bo Qiao, and Shilin He, *Microsoft Research*; Pochian Lee, Jeffrey Sun, Feng Gao, and Li Yang, *Microsoft Azure*; Qingwei Lin, *Microsoft Research*; Saravanakumar Rajmohan, *Microsoft 365*; Zhangwei Xu, *Microsoft Azure*; Dongmei Zhang, *Microsoft Research*

SIT, Fido!: Training Machine Learning Algorithms

- Habitat: A Runtime-Based Computational Performance Predictor for Deep Neural Network Training** 503
Geoffrey X. Yu, *University of Toronto/Vector Institute*; Yubo Gao, *University of Toronto*; Pavel Golikov and Gennady Pekhimenko, *University of Toronto/Vector Institute*
- Zico: Efficient GPU Memory Sharing for Concurrent DNN Training** 523
Gangmuk Lim, *UNIST*; Jeongseob Ahn, *Ajou University*; Wencong Xiao, *Alibaba Group*; Youngjin Kwon, *KAIST*; Myeongjae Jeon, *UNIST*
- Refurbish Your Training Data: Reusing Partially Augmented Samples for Faster Deep Neural Network Training** ... 537
Gyewon Lee, *Seoul National University and FriendliAI*; Irene Lee, *Georgia Institute of Technology*; Hyeonmin Ha, Kyungeun Lee, and Hwarim Hyun, *Seoul National University*; Ahnjae Shin and Byung-Gon Chun, *Seoul National University and FriendliAI*
- ZeRO-Offload: Democratizing Billion-Scale Model Training** 551
Jie Ren, *UC Merced*; Samyam Rajbhandari, Reza Yazdani Aminabadi, and Olatunji Ruwase, *Microsoft*; Shuangyan Yang, *UC Merced*; Minjia Zhang, *Microsoft*; Dong Li, *UC Merced*; Yuxiong He, *Microsoft*

I Can Smell That Fluffy Was Here: Networks

- Hashing Linearity Enables Relative Path Control in Data Centers** 565
Zhehui Zhang, *University of California, Los Angeles*; Haiyang Zheng, Jiayao Hu, Xiangning Yu, Chenchen Qi, Xuemei Shi, and Guohui Wang, *Alibaba Group*
- Live in the Express Lane** 581
Patrick Jahnke, *TU Darmstadt and SAP*; Vincent Riesop, *SAP*; Pierre-Louis Roman and Pavel Chuprikov, *Università della Svizzera italiana*; Patrick Eugster, *Università della Svizzera italiana, TU Darmstadt, and Purdue University*
- Understanding Precision Time Protocol in Today's Wi-Fi Networks: A Measurement Study** 597
Paizhuo Chen and Zhice Yang, *ShanghaiTech University*
- AUTO: Adaptive Congestion Control Based on Multi-Objective Reinforcement Learning for the Satellite-Ground Integrated Network** 611
Xu Li, Feilong Tang, and Jiacheng Liu, *Shanghai Jiao Tong University*; Laurence T. Yang, *St. Francis Xavier University*; Luoyi Fu and Long Chen, *Shanghai Jiao Tong University*
- Hey, Lumi! Using Natural Language for Intent-Based Network Management** 625
Arthur S. Jacobs, Ricardo J. Pfitscher, and Rafael H. Ribeiro, *Federal University of Rio Grande do Sul (UFRGS)*; Ronaldo A. Ferreira, *UFMS*; Lisandro Z. Granville, *Federal University of Rio Grande do Sul (UFRGS)*; Walter Willinger, *NIKSUN, Inc.*; Sanjay G. Rao, *Purdue University*

I Buried That Bone Here Somewhere: Storage

- Boosting Full-Node Repair in Erasure-Coded Storage.** 641
Shiyao Lin, Guowen Gong, and Zhirong Shen, *Xiamen University*; Patrick P. C. Lee, *The Chinese University of Hong Kong*; Jiwu Shu, *Xiamen University and Tsinghua University*
- KVIMR: Key-Value Store Aware Data Management Middleware for Interlaced Magnetic Recording Based Hard Disk Drive** 657
Yuhong Liang, Tsun-Yu Yang, and Ming-Chang Yang, *The Chinese University of Hong Kong*
- Differentiated Key-Value Storage Management for Balanced I/O Performance** 673
Yongkun Li and Zhen Liu, *University of Science and Technology of China*; Patrick P. C. Lee, *The Chinese University of Hong Kong*; Jiayu Wu and Yinlong Xu, *Anhui Province Key Laboratory of High Performance Computing, University of Science and Technology of China*; Yi Wu, Liu Tang, Qi Liu, and Qiu Cui, *PingCAP*
- ZNS: Avoiding the Block Interface Tax for Flash-based SSDs** 689
Matias Björling, *Western Digital*; Abutalib Aghayev, *The Pennsylvania State University*; Hans Holmberg, Aravind Ramesh, and Damien Le Moal, *Western Digital*; Gregory R. Ganger and George Amvrosiadis, *Carnegie Mellon University*
- MapperX: Adaptive Metadata Maintenance for Fast Crash Recovery of DM-Cache Based Hybrid Storage Devices.** ..705
Lujia Yin, *NUDT*; Li Wang, *Didi Chuxing*; Yiming Zhang, *NiceX Lab, NUDT*; Yuxing Peng, *NUDT*

Friday, July 16

My Tail Never Has Any Latency: OS & Hardware

- Exploring the Design Space of Page Management for Multi-Tiered Memory Systems.** 715
Jonghyeon Kim, Wonkyo Choe, and Jeongseob Ahn, *Ajou University*
- A Fast and Flexible Hardware-based Virtualization Mechanism for Computational Storage Devices.** 729
Dongup Kwon, Dongryeong Kim, Junehyuk Boo, Wonsik Lee, and Jangwoo Kim, *Seoul National University*
- Fair Scheduling for AVX2 and AVX-512 Workloads** 745
Mathias Gottschlag, Philipp Machauer, Yussuf Khalil, and Frank Bellosa, *Karlsruhe Institute of Technology*
- SKQ: Event Scheduling for Optimizing Tail Latency in a Traditional OS Kernel** 759
Siyao Zhao, Haoyu Gu, and Ali José Mashtizadeh, *University of Waterloo*
- A Linux Kernel Implementation of the Homa Transport Protocol.** 773
John Ousterhout, *Stanford University*

Friends Fur-Ever: Persistent Memory and In-Memory Computing

- Ayudante: A Deep Reinforcement Learning Approach to Assist Persistent Memory Programming** 789
Hanxian Huang, Zixuan Wang, Juno Kim, Steven Swanson, and Jishen Zhao, *University of California, San Diego*
- Tips: Making Volatile Index Structures Persistent with DRAM-NVMM Tiering.** 805
R. Madhava Krishnan, Wook-Hee Kim, Xinwei Fu, and Sumit Kumar Monga, *Virginia Tech*; Hee Won Lee, *Samsung Electronics*; Minsung Jang, *Peraton Labs*; Ajit Mathew and Changwoo Min, *Virginia Tech*
- Improving Performance of Flash Based Key-Value Stores Using Storage Class Memory as a Volatile Memory Extension** 821
Hiwot Tadese Kassa, *University of Michigan*; Jason Akers, Mrinmoy Ghosh, and Zhichao Cao, *Facebook Inc.*; Vaibhav Gogte and Ronald Dreslinski, *University of Michigan*
- First Responder: Persistent Memory Simultaneously as High Performance Buffer Cache and Storage** 839
Hyunsub Song, Shean Kim, J. Hyun Kim, Ethan JH Park, and Sam H. Noh, *UNIST*
- A Case Study of Processing-in-Memory in off-the-Shelf Systems** 855
Joel Nider, Craig Mustard, Andrada Zoltan, John Ramsden, Larry Liu, Jacob Grossbard, and Mohammad Dashti, *University of British Columbia*; Romaric Jodin, Alexandre Ghiti, and Jordi Chauzi, *UPMEM SAS*; Alexandra Fedorova, *University of British Columbia*

Time to File the Claws: Files

- XFUSE: An Infrastructure for Running Filesystem Services in User Space** 863
Qianbo Huai, Windsor Hsu, Jiwei Lu, Hao Liang, Haobo Xu, and Wei Chen, *Alibaba Group*
- MAX: A Multicore-Accelerated File System for Flash Storage** 877
Xiaojian Liao, Youyou Lu, Erci Xu, and Jiwu Shu, *Department of Computer Science and Technology, Tsinghua University, and Beijing National Research Center for Information Science and Technology (BNRist)*
- Z-Journal: Scalable Per-Core Journaling** 893
Jongseok Kim and Cassiano Campes, *Sungkyunkwan University*; Joo-Young Hwang, *Samsung Electronics Co., Ltd.*; Jinkyu Jeong and Euseong Seo, *Sungkyunkwan University*
- LODIC: Logical Distributed Counting for Scalable File Access** 907
Jeoungahn Park, *KAIST*; Taeho Hwang, *Hanyang University*; Jongmoo Choi, *Dankook University*; Changwoo Min, *Virginia Tech*; Youjip Won, *KAIST*

But You Played with Me Yesterday: Serverless Computing and Consistency

- UNISTORE: A fault-tolerant marriage of causal and strong consistency** 923
Manuel Bravo, Alexey Gotsman, and Borja de Régil, *IMDEA Software Institute*; Hengfeng Wei, *Nanjing University*
- Optimistic Concurrency Control for Real-world Go Programs** 939
Zhizhou Zhang, *University of California, Santa Barbara*; Milind Chabbi and Adam Welc, *Uber Technologies*; Timothy Sherwood, *University of California, Santa Barbara*
- Faastlane: Accelerating Function-as-a-Service Workflows** 957
Swaroop Kotni, Ajay Nayak, Vinod Ganapathy, and Arkaprava Basu, *Indian Institute of Science*
- SONIC: Application-aware Data Passing for Chained Serverless Applications** 973
Ashraf Mahgoub, *Purdue University*; Karthick Shankar, *Carnegie Mellon University*; Subrata Mitra, *Adobe Research*; Ana Klimovic, *ETH Zurich*; Somali Chaterji and Saurabh Bagchi, *Purdue University*