

21st USENIX Symposium on Networked Systems Design and Implementation (NSDI'24)

Santa Clara, California, USA
16-18 April 2024

Volume 1 of 3

ISBN: 978-1-7138-9652-4

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© (2024) by Usenix Association
All rights reserved.

Printed with permission by Curran Associates, Inc. (2024)

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

21st USENIX Symposium on Networked Systems Design and Implementation (NSDI '24)

April 16–18, 2024
Santa Clara, CA, USA

Tuesday, April 16

Clouds but Faster

- Horus: Granular In-Network Task Scheduler for Cloud Datacenters** 1
Parham Yassini, *Simon Fraser University*; Khaled Diab, *Hewlett Packard Labs*; Saeed Zangeneh and Mohamed Hefeeda, *Simon Fraser University*
- Fast Vector Query Processing for Large Datasets Beyond GPU Memory with Reordered Pipelining** 23
Zili Zhang, Fangyue Liu, Gang Huang, Xuanzhe Liu, and Xin Jin, *School of Computer Science, Peking University*
- LoLKV: The Logless, Linearizable, RDMA-based Key-Value Storage System** 41
Ahmed Alquraan and Sreeharsha Udayashankar, *University of Waterloo*; Virendra Marathe, *Oracle Labs*; Bernard Wong and Samer Al-Kiswany, *University of Waterloo*
- Making Kernel Bypass Practical for the Cloud with Junction** 55
Joshua Fried and Gohar Irfan Chaudhry, *MIT CSAIL*; Enrique Saurez, Esha Choukse, and Íñigo Goiri, *Azure Research – Systems*; Sameh Elnikety, *Microsoft Research*; Rodrigo Fonseca, *Azure Research – Systems*; Adam Belay, *MIT CSAIL*

Scheduling the Network

- Sifter: An Inversion-Free and Large-Capacity Programmable Packet Scheduler** 75
Peixuan Gao, Anthony Dalleggio, Jiajin Liu, and Chen Peng, *New York University*; Yang Xu, *Fudan University*; H. Jonathan Chao, *New York University*
- Flow Scheduling with Imprecise Knowledge** 95
Wenxin Li, Xin He, Yuan Liu, and Keqiu Li, *Tianjin University*; Kai Chen, *Hong Kong University of Science and Technology and University of Science and Technology of China*; Zhao Ge and Zewei Guan, *Tianjin University*; Heng Qi, *Dalian University of Technology*; Song Zhang, *Tianjin University*; Guyue Liu, *New York University Shanghai*
- Pudica: Toward Near-Zero Queuing Delay in Congestion Control for Cloud Gaming.** 113
Shibo Wang, *Xi'an Jiaotong University and Tencent Inc.*; Shusen Yang, *Xi'an Jiaotong University*; Xiao Kong, Chenglei Wu, and Longwei Jiang, *Tencent Inc.*; Chenren Xu, *Peking University*; Cong Zhao, *Xi'an Jiaotong University*; Xuesong Yang, *Bonree*; Jianjun Xiao and Xin Liu, *Tencent Inc.*; Changxi Zheng, *Pixel Lab, Tencent America, and Columbia University*; Jing Wang and Honghao Liu, *Tencent Inc.*
- Revisiting Congestion Control for Lossless Ethernet** 131
Yiran Zhang, *Tsinghua University and Beijing University of Posts and Telecommunications*; Qingkai Meng, *Tsinghua University and Beihang University*; Chaolei Hu and Fengyuan Ren, *Tsinghua University*

Serverless

- Autothrottle: A Practical Bi-Level Approach to Resource Management for SLO-Targeted Microservices** 149
Zibo Wang, *University of Science and Technology of China and Microsoft Research*; Pinghe Li, *ETH Zurich*; Chieh-Jan Mike Liang, *Microsoft Research*; Feng Wu, *University of Science and Technology of China*; Francis Y. Yan, *Microsoft Research*
- Jolteon: Unleashing the Promise of Serverless for Serverless Workflows** 167
Zili Zhang, Chao Jin, and Xin Jin, *School of Computer Science, Peking University*
- Can't Be Late: Optimizing Spot Instance Savings under Deadlines** 185
Zhanghao Wu, Wei-Lin Chiang, Ziming Mao, and Zongheng Yang, *University of California, Berkeley*; Eric Friedman and Scott Shenker, *University of California, Berkeley, and ICSI*; Ion Stoica, *University of California, Berkeley*
- Towards Intelligent Automobile Cockpit via A New Container Architecture** 205
Lin Jiang and Feiyu Zhang, *Xi'an Yunzhi Technology*; Jiang Ming, *Tulane University*

| | |
|---|------------|
| MuCache: a General Framework for Caching in Microservice Graphs | 221 |
| Haoran Zhang, Konstantinos Kallas, Spyros Pavlatos, Rajeev Alur, Sebastian Angel, and Vincent Liu, <i>University of Pennsylvania</i> | |
| Network Protocols | |
| A large-scale deployment of DCTCP | 239 |
| Abhishek Dhamija and Balasubramanian Madhavan, <i>Meta</i> ; Hechao Li, <i>Netflix</i> ; Jie Meng, Shrikrishna Khare, and Madhavi Rao, <i>Meta</i> ; Lawrence Brakmo; Neil Spring, Prashanth Kannan, and Srikanth Sundaresan, <i>Meta</i> ; Soudeh Ghorbani, <i>Meta and Johns Hopkins University</i> | |
| TECC: Towards Efficient QUIC Tunneling via Collaborative Transmission Control | 253 |
| Jiaxing Zhang, <i>Alibaba Group and University of Chinese Academy of Sciences</i> ; Furong Yang, <i>Alibaba Group</i> ; Ting Liu, <i>Alibaba Group and University of Chinese Academy of Sciences</i> ; Qinghua Wu, <i>University of Chinese Academy of Sciences and Purple Mountain Laboratories, China</i> ; Wu Zhao, Yuanbo Zhang, Wentao Chen, Yanmei Liu, Hongyu Guo, and Yunfei Ma, <i>Alibaba Group</i> ; Zhenyu Li, <i>University of Chinese Academy of Sciences and Purple Mountain Laboratories, China</i> | |
| iStack: A General and Stateful Name-based Protocol Stack for Named Data Networking | 267 |
| Tianlong Li, Tian Song, and Yating Yang, <i>Beijing Institute of Technology</i> | |
| Cloudcast: High-Throughput, Cost-Aware Overlay Multicast in the Cloud | 281 |
| Sarah Wooders and Shu Liu, <i>UC Berkeley</i> ; Paras Jain, <i>Genmo AI</i> ; Xiangxi Mo and Joseph Gonzalez, <i>UC Berkeley</i> ; Vincent Liu, <i>University of Pennsylvania</i> ; Ion Stoica, <i>UC Berkeley</i> | |
| Understanding Routable PCIe Performance for Composable Infrastructures | 297 |
| Wentao Hou, <i>University of Wisconsin-Madison</i> ; Jie Zhang and Zeke Wang, <i>Zhejiang University</i> ; Ming Liu, <i>University of Wisconsin-Madison</i> | |
| Distributed Systems: Part 1 | |
| Alea-BFT: Practical Asynchronous Byzantine Fault Tolerance | 313 |
| Diogo S. Antunes, Afonso N. Oliveira, André Breda, Matheus Guilherme Franco, Henrique Moniz, and Rodrigo Rodrigues, <i>Instituto Superior Técnico (ULisboa) and INESC-ID</i> | |
| Harmony: A Congestion-free Datacenter Architecture | 329 |
| Saksham Agarwal, Qizhe Cai, Rachit Agarwal, and David Shmoys, <i>Cornell University</i> ; Amin Vahdat, <i>Google</i> | |
| SwiftPaxos: Fast Geo-Replicated State Machines | 345 |
| Fedor Ryabinin, <i>IMDEA Software Institute and Universidad Politécnica de Madrid</i> ; Alexey Gotsman, <i>IMDEA Software Institute</i> ; Pierre Sutra, <i>Télécom SudParis and INRIA</i> | |
| The Bedrock of Byzantine Fault Tolerance: A Unified Platform for BFT Protocols Analysis, Implementation, and Experimentation | 371 |
| Mohammad Javad Amiri, <i>Stony Brook University</i> ; Chenyuan Wu, <i>University of Pennsylvania</i> ; Divyakant Agrawal and Amr El Abbadi, <i>UC Santa Barbara</i> ; Boon Thau Loo, <i>University of Pennsylvania</i> ; Mohammad Sadoghi, <i>UC Davis</i> | |
| DINT: Fast In-Kernel Distributed Transactions with eBPF | 401 |
| Yang Zhou, <i>Harvard University</i> ; Xingyu Xiang, <i>Peking University</i> ; Matthew Kiley, <i>Harvard University</i> ; Sowmya Dharanipragada, <i>Cornell University</i> ; Minlan Yu, <i>Harvard University</i> | |
| Programming the Network: Part 1 | |
| Brain-on-Switch: Towards Advanced Intelligent Network Data Plane via NN-Driven Traffic Analysis at Line-Speed | 419 |
| Jinzhu Yan and Haotian Xu, <i>Tsinghua University</i> Zhuotao Liu, Qi Li, Ke Xu, Mingwei Xu, and Jianping Wu, <i>Tsinghua University and Zhongguancun Laboratory</i> | |
| The Eternal Tussle: Exploring the Role of Centralization in IPFS | 441 |
| Yiluo Wei, <i>Hong Kong University of Science & Technology (GZ)</i> ; Dennis Trautwein and Yiannis Psaras, <i>Protocol Labs</i> ; Ignacio Castro, <i>Queen Mary University of London</i> ; Will Scott, <i>Protocol Labs</i> ; Aravindh Raman, <i>Brave Software</i> ; Gareth Tyson, <i>Hong Kong University of Science & Technology (GZ)</i> | |
| BBQ: A Fast and Scalable Integer Priority Queue for Hardware Packet Scheduling | 455 |
| Nirav Atre, Hugo Sadok, and Justine Sherry, <i>Carnegie Mellon University</i> | |

| | |
|--|------------|
| Sirius: Composing Network Function Chains into P4-Capable Edge Gateways | 477 |
| Jiaqi Gao, Jiamin Cao, Yifan Li, Mengqi Liu, Ming Tang, Dennis Cai, and Ennan Zhai, <i>Alibaba Cloud</i> | |
| Empower Programmable Pipeline for Advanced Stateful Packet Processing | 491 |
| Yong Feng and Zhikang Chen, <i>Tsinghua University</i> ; Haoyu Song, <i>Futurewei Technologies</i> ; Yinchao Zhang, Hanyi Zhou, Ruoyu Sun, Wenkuo Dong, Peng Lu, Shuxin Liu, and Chuwen Zhang, <i>Tsinghua University</i> ; Yang Xu, <i>Fudan University</i> ; Bin Liu, <i>Tsinghua University</i> | |

Video

| | |
|---|------------|
| GRACE: Loss-Resilient Real-Time Video through Neural Codecs | 509 |
| Yihua Cheng, Ziyi Zhang, Hanchen Li, Anton Arapin, and Yue Zhang, <i>The University of Chicago</i> ; Qizheng Zhang, <i>Stanford University</i> ; Yuhan Liu, Kuntai Du, and Xu Zhang, <i>The University of Chicago</i> ; Francis Y. Yan, <i>Microsoft</i> ; Amrita Mazumdar, <i>NVIDIA</i> ; Nick Feamster and Junchen Jiang, <i>The University of Chicago</i> | |
| LiFeR: Unleash Learned Codecs in Video Streaming with Loose Frame Referencing | 533 |
| Bo Chen, <i>University of Illinois at Urbana-Champaign</i> ; Zhisheng Yan, <i>George Mason University</i> ; Yinjie Zhang, Zhe Yang, and Klara Nahrstedt, <i>University of Illinois at Urbana-Champaign</i> | |
| MadEye: Boosting Live Video Analytics Accuracy with Adaptive Camera Configurations | 549 |
| Mike Wong and Murali Ramanujam, <i>Princeton University</i> ; Guha Balakrishnan, <i>Rice University</i> ; Ravi Netravali, <i>Princeton University</i> | |
| Gemino: Practical and Robust Neural Compression for Video Conferencing | 569 |
| Vibhaalakshmi Sivaraman, Pantea Karimi, Vedantha Venkatapathy, and Mehrdad Khani, <i>Massachusetts Institute of Technology</i> ; Sadjad Fouladi, <i>Microsoft Research</i> ; Mohammad Alizadeh, Frédo Durand, and Vivienne Sze, <i>Massachusetts Institute of Technology</i> | |
| ARTEMIS: Adaptive Bitrate Ladder Optimization for Live Video Streaming | 591 |
| Farzad Tashtarian, <i>Christian Doppler Laboratory ATHENA, Alpen-Adria Universität Klagenfurt</i> ; Abdelhak Bentaleb, <i>Concordia University</i> ; Hadi Amirpour, <i>Christian Doppler Laboratory ATHENA, Alpen-Adria Universität Klagenfurt</i> ; Sergey Gorinsky, <i>IMDEA Networks Institute</i> ; Junchen Jiang, <i>University of Chicago</i> ; Hermann Hellwagner and Christian Timmerer, <i>Christian Doppler Laboratory ATHENA, Alpen-Adria Universität Klagenfurt</i> | |

Sharing the Network

| | |
|--|------------|
| Credence: Augmenting Datacenter Switch Buffer Sharing with ML Predictions | 613 |
| Vamsi Addanki, Maciej Pacut, and Stefan Schmid, <i>TU Berlin</i> | |
| Seer: Enabling Future-Aware Online Caching in Networked Systems | 635 |
| Jason Lei and Vishal Shrivastav, <i>Purdue University</i> | |
| Reverie: Low Pass Filter-Based Switch Buffer Sharing for Datacenters with RDMA and TCP Traffic | 651 |
| Vamsi Addanki, <i>TU Berlin</i> ; Wei Bai, <i>Microsoft Research</i> ; Stefan Schmid, <i>TU Berlin</i> ; Maria Apostolaki, <i>Princeton University</i> | |
| Precise Data Center Traffic Engineering with Constrained Hardware Resources | 669 |
| Shawn Shuoshuo Chen, <i>Carnegie Mellon University</i> ; Keqiang He, <i>Airbnb</i> ; Rui Wang, <i>Google</i> ; Srinivasan Seshan and Peter Steenkiste, <i>Carnegie Mellon University</i> | |
| Multitenant In-Network Acceleration with SwitchVM | 691 |
| Sajy Khashab, Alon Rashelbach, and Mark Silberstein, <i>Technion</i> | |

Wednesday, April 17

ML at Scale

| | |
|--|------------|
| Characterization of Large Language Model Development in the Datacenter | 709 |
| Qinghao Hu, <i>Shanghai AI Laboratory and S-Lab, Nanyang Technological University</i> ; Zhisheng Ye, <i>Shanghai AI Laboratory and Peking University</i> ; Zerui Wang, <i>Shanghai AI Laboratory and Shanghai Jiao Tong University</i> ; Guoteng Wang, <i>Shanghai AI Laboratory</i> ; Meng Zhang and Qiaoling Chen, <i>Shanghai AI Laboratory and S-Lab, Nanyang Technological University</i> ; Peng Sun, <i>Shanghai AI Laboratory and SenseTime Research</i> ; Dahua Lin, <i>Shanghai AI Laboratory and CUHK</i> ; Xiaolin Wang and Yingwei Luo, <i>Peking University</i> ; Yonggang Wen and Tianwei Zhang, <i>Nanyang Technological University</i> | |

QuickUpdate: a Real-Time Personalization System for Large-Scale Recommendation Models 731
Kiran Kumar Matam, Hani Ramezani, Fan Wang, Zeliang Chen, Yue Dong, Maomao Ding, Zhiwei Zhao, Zhengyu Zhang, Ellie Wen, and Assaf Eisenman, *Meta, Inc.*

MegaScale: Scaling Large Language Model Training to More Than 10,000 GPUs 745
Ziheng Jiang and Haibin Lin, *ByteDance*; Yinmin Zhong, *Peking University*; Qi Huang, Yangrui Chen, Zhi Zhang, Yanghua Peng, Xiang Li, Cong Xie, Shibiao Nong, Yulu Jia, Sun He, Hongmin Chen, Zhihao Bai, Qi Hou, Shipeng Yan, Ding Zhou, Yiyao Sheng, Zhuo Jiang, Haohan Xu, Haoran Wei, Zhang Zhang, Pengfei Nie, Leqi Zou, Sida Zhao, Liang Xiang, Zherui Liu, Zhe Li, Xiaoying Jia, and Jianxi Ye, *ByteDance*; Xin Jin, *Peking University*; Xin Liu, *ByteDance*

Resiliency at Scale: Managing Google’s TPUv4 Machine Learning Supercomputer 761
Yazhou Zu, Alireza Ghaffarkhah, Hoang-Vu Dang, Brian Towles, Steven Hand, Safeen Huda, Adekunle Bello, Alexander Kolbasov, Arash Rezaei, Dayou Du, Steve Lacy, Hang Wang, Aaron Wisner, Chris Lewis, and Henri Bahini, *Google*

Satellites and Things

NN-Defined Modulator: Reconfigurable and Portable Software Modulator on IoT Gateways 775
Jiazhao Wang and Wenchao Jiang, *Singapore University of Technology and Design*; Ruofeng Liu, *University of Minnesota*; Bin Hu, *University of Southern California*; Demin Gao, *Nanjing Forestry University*; Shuai Wang, *Southeast University*

Democratizing Direct-to-Cell Low Earth Orbit Satellite Networks 791
Lixin Liu, *Tsinghua University*; Yuanjie Li and Hewu Li, *Tsinghua University and Zhongguancun Laboratory*; Jiabo Yang, Wei Liu, Jingyi Lan, Yufeng Wang, and Jiarui Li, *Tsinghua University*; Jianping Wu, Qian Wu, Jun Liu, and Zeqi Lai, *Tsinghua University and Zhongguancun Laboratory*

Known Knowns and Unknowns: Near-realtime Earth Observation Via Query Bifurcation in Serval 809
Bill Tao, Om Chabra, Ishani Janveja, Indranil Gupta, and Deepak Vasisht, *University of Illinois Urbana-Champaign*

Spectrumize: Spectrum-efficient Satellite Networks for the Internet of Things 825
Vaibhav Singh, Tusher Chakraborty, and Suraj Jog, *Microsoft Research*; Om Chabra and Deepak Vasisht, *UIUC*; Ranveer Chandra, *Microsoft Research*

Wide-Area and Edge

Application-Level Service Assurance with 5G RAN Slicing 841
Arjun Balasingam, *MIT CSAIL*; Manikanta Kotaru and Paramvir Bahl, *Microsoft*

CHISEL: An optical slice of the wide-area network 859
Abhishek Vijaya Kumar, *Cornell University*; Bill Owens, *NYSErnet*; Nikolaj Bjørner, Binbin Guan, Yawei Yin, and Paramvir Bahl, *Microsoft*; Rachee Singh, *Cornell University*

LuoShen: A Hyper-Converged Programmable Gateway for Multi-Tenant Multi-Service Edge Clouds 877
Tian Pan, Kun Liu, Xionglie Wei, Yisong Qiao, Jun Hu, Zhiguo Li, Jun Liang, Tiesheng Cheng, Wenqiang Su, Jie Lu, Yuke Hong, Zhengzhong Wang, Zhi Xu, Chongjing Dai, Peiqiao Wang, Xuetao Jia, Jianyuan Lu, Enge Song, and Jun Zeng, *Alibaba Cloud*; Biao Lyu, *Zhejiang University and Alibaba Cloud*; Ennan Zhai, *Alibaba Cloud*; Jiao Zhang and Tao Huang, *Purple Mountain Laboratories*; Dennis Cai, *Alibaba Cloud*; Shunmin Zhu, *Tsinghua University and Alibaba Cloud*

Sprinter: Speeding Up High-Fidelity Crawling of the Modern Web 893
Ayush Goel and Jingyuan Zhu, *University of Michigan*; Ravi Netravali, *Princeton University*; Harsha V. Madhyastha, *University of Southern California*

Hairpin: Rethinking Packet Loss Recovery in Edge-based Interactive Video Streaming 907
Zili Meng, *Tsinghua University, Hong Kong University of Science and Technology, and Tencent*; Xiao Kong and Jing Chen, *Tsinghua University and Tencent*; Bo Wang and Mingwei Xu, *Tsinghua University*; Rui Han and Honghao Liu, *Tencent*; Venkat Arun, *UT Austin*; Hongxin Hu, *University at Buffalo, SUNY*; Xue Wei, *Tencent*

Verification

Finding Adversarial Inputs for Heuristics using Multi-level Optimization 927
Pooria Namyar, *Microsoft and University of Southern California*; Behnaz Arzani and Ryan Beckett, *Microsoft*; Santiago Segarra, *Microsoft and Rice University*; Himanshu Raj and Umesh Krishnaswamy, *Microsoft*; Ramesh Govindan, *University of Southern California*; Srikanth Kandula, *Microsoft*

| | |
|--|-------------|
| Towards provably performant congestion control. | 951 |
| Anup Agarwal, <i>Carnegie Mellon University</i> ; Venkat Arun, <i>University of Texas at Austin</i> ; Devdeep Ray, Ruben Martins, and Srinivasan Seshan, <i>Carnegie Mellon University</i> | |
| EPVerifier: Accelerating Update Storms Verification with Edge-Predicate | 979 |
| Chenyang Zhao, Yuebin Guo, Jingyu Wang, Qi Qi, Zirui Zhuang, Haifeng Sun, and Lingqi Guo, <i>State Key Laboratory of Networking and Switching Technology, Beijing University of Posts and Telecommunications</i> ; Yuming Xie, <i>Huawei Technologies Co., Ltd</i> ; Jianxin Liao, <i>State Key Laboratory of Networking and Switching Technology, Beijing University of Posts and Telecommunications</i> | |
| Netcastle: Network Infrastructure Testing At Scale | 993 |
| Rob Sherwood, <i>NetDebug.com</i> ; Jinghao Shi, Ying Zhang, Neil Spring, Srikanth Sundaresan, Jasmeet Bagga, Prathyusha Peddi, Vineela Kukkadapu, Rashmi Shrivastava, Manikantan KR, Pavan Patil, Srikrishna Gopu, Varun Varadan, Ethan Shi, Hany Morsy, Yuting Bu, Renjie Yang, Rasmus Jönsson, Wei Zhang, Jesus Jussepén Arredondo, and Diana Saha, <i>Meta Platforms Inc.</i> ; Sean Choi, <i>Santa Clara University</i> | |
| MESSI: Behavioral Testing of BGP Implementations | 1009 |
| Rathin Singha and Rajdeep Mondal, <i>University of California Los Angeles</i> ; Ryan Beckett, <i>Microsoft</i> ; Siva Kesava Reddy Kakarla, <i>Microsoft Research</i> ; Todd Millstein and George Varghese, <i>University of California Los Angeles</i> | |
| Networking at Scale | |
| A High-Performance Design, Implementation, Deployment, and Evaluation of The Slim Fly Network. | 1025 |
| Nils Blach and Maciej Besta, <i>ETH Zürich</i> ; Daniele De Sensi, <i>ETH Zürich and Sapienza University of Rome</i> ; Jens Domke, <i>RIKEN Center for Computational Science (R-CCS)</i> ; Hussein Harake, <i>Swiss National Supercomputing Centre (CSCS)</i> ; Shigang Li, <i>ETH Zürich and BUPT, Beijing</i> ; Patrick Iff, <i>ETH Zürich</i> ; Marek Konieczny, <i>AGH-UST</i> ; Kartik Lakhota, <i>Intel Labs</i> ; Ales Kubicek and Marcel Ferrari, <i>ETH Zürich</i> ; Fabrizio Petrini, <i>Intel Labs</i> ; Torsten Hoefler, <i>ETH Zürich</i> | |
| Crescent: Emulating Heterogeneous Production Network at Scale. | 1045 |
| Zhaoyu Gao and Anubhavnidhi Abhashkumar, <i>ByteDance</i> ; Zhen Sun, <i>Cornell University</i> ; Weirong Jiang and Yi Wang, <i>ByteDance</i> | |
| Reasoning about Network Traffic Load Property at Production Scale | 1063 |
| Ruihan Li, <i>Peking University and Alibaba Cloud</i> ; Fangdan Ye, Yifei Yuan, Ruizhen Yang, Bingchuan Tian, Tianchen Guo, Hao Wu, Xiaobo Zhu, Zhongyu Guan, Qing Ma, and Xianlong Zeng, <i>Alibaba Cloud</i> ; Chenren Xu, <i>Peking University</i> ; Dennis Cai and Ennan Zhai, <i>Alibaba Cloud</i> | |
| POSEIDON: A Consolidated Virtual Network Controller that Manages Millions of Tenants via Config Tree | 1083 |
| Biao Lyu, <i>Zhejiang University and Alibaba Cloud</i> ; Enge Song, Tian Pan, Jianyuan Lu, Shize Zhang, Xiaoqing Sun, Lei Gao, Chenxiao Wang, Han Xiao, Yong Pan, Xiuheng Chen, Yandong Duan, Weisheng Wang, Jinpeng Long, Yanfeng Wang, Kunpeng Zhou, and Zhigang Zong, <i>Alibaba Cloud</i> ; Xing Li, <i>Zhejiang University and Alibaba Cloud</i> ; Guangwang Li and Pengyu Zhang, <i>Alibaba Cloud</i> ; Peng Cheng and Jiming Chen, <i>Zhejiang University</i> ; Shunmin Zhu, <i>Tsinghua University and Alibaba Cloud</i> | |
| OPPerTune: Post-Deployment Configuration Tuning of Services Made Easy. | 1101 |
| Gagan Somashekar, <i>Stony Brook University</i> ; Karan Tandon and Anush Kini, <i>Microsoft Research</i> ; Chieh-Chun Chang and Petr Husak, <i>Microsoft</i> ; Ranjita Bhagwan, <i>Google</i> ; Mayukh Das, <i>Microsoft365 Research</i> ; Anshul Gandhi, <i>Stony Brook University</i> ; Nagarajan Natarajan, <i>Microsoft Research</i> | |
| ML but Faster | |
| Parcae: Proactive, Liveput-Optimized DNN Training on Preemptible Instances | 1121 |
| Jiangfei Duan, <i>The Chinese University of Hong Kong</i> ; Ziang Song, <i>ByteDance</i> ; Xupeng Miao and Xiaoli Xi, <i>Carnegie Mellon University</i> ; Dahua Lin, <i>The Chinese University of Hong Kong</i> ; Harry Xu, <i>University of California, Los Angeles</i> ; Minjia Zhang, <i>Microsoft</i> ; Zhihao Jia, <i>Carnegie Mellon University</i> | |
| Accelerating Neural Recommendation Training with Embedding Scheduling | 1141 |
| Chaoliang Zeng, Xudong Liao, Xiaodian Cheng, Han Tian, Xinchun Wan, Hao Wang, and Kai Chen, <i>iSING Lab, Hong Kong University of Science and Technology</i> | |
| DISTMM: Accelerating Distributed Multimodal Model Training | 1157 |
| Jun Huang, <i>The Ohio State University</i> ; Zhen Zhang, <i>Amazon Web Services</i> ; Shuai Zheng, <i>Boson AI</i> ; Feng Qin, <i>The Ohio State University</i> ; Yida Wang, <i>Amazon Web Services</i> | |

Approximate Caching for Efficiently Serving Text-to-Image Diffusion Models1173
Shubham Agarwal and Subrata Mitra, *Adobe Research*; Sarthak Chakraborty, *UIUC*; Srikrishna Karanam,
Koyel Mukherjee, and Shiv Kumar Saini, *Adobe Research*

THC: Accelerating Distributed Deep Learning Using Tensor Homomorphic Compression 1191
Minghao Li, *Harvard University*; Ran Ben Basat, *University College London*; Shay Vargaftik, *VMware Research*;
ChonLam Lao, Kevin Xu, Michael Mitzenmacher, and Minlan Yu, *Harvard University*

Distributed Systems: Part 2

Accelerating Skewed Workloads With Performance Multipliers in the TurboDB Distributed Database 1213
Jennifer Lam, Jeffrey Helt, and Wyatt Lloyd, *Princeton University*; Haonan Lu, *University at Buffalo*

SIEVE is Simpler than LRU: an Efficient Turn-Key Eviction Algorithm for Web Caches 1229
Yazhuo Zhang, *Emory University*; Juncheng Yang, *Carnegie Mellon University*; Yao Yue, *Pelikan Foundation*;
Ymir Vigfusson, *Emory University and Keystrike*; K.V. Rashmi, *Carnegie Mellon University*

Harvesting Idle Memory for Application-Managed Soft State with Midas 1247
Yifan Qiao, *UCLA*; Zhenyuan Ruan, *MIT CSAIL*; Haoran Ma, *UCLA*; Adam Belay, *MIT CSAIL*; Miryung Kim
and Harry Xu, *UCLA*

Efficient Exposure of Partial Failure Bugs in Distributed Systems with Inferred Abstract States 1267
Haoze Wu and Jia Pan, *Johns Hopkins University*; Peng Huang, *University of Michigan*

Load is not what you should balance: Introducing Prequal 1285
Bartek Wydrowski, *Google Research*; Robert Kleinberg, *Google Research and Cornell*; Stephen M. Rumble, *Google (YouTube)*;
Aaron Archer, *Google Research*

Wireless Hardware

Orthcatter: High-throughput In-band OFDM Backscatter with Over-the-Air Code Division 1301
Caihui Du and Jihong Yu, *Beijing Institute of Technology*; Rongrong Zhang, *Capital Normal University*;
Ju Ren, *Tsinghua University*; Jianping An, *Beijing Institute of Technology*

EdgeRIC: Empowering Realtime Intelligent Optimization and Control in NextG Cellular Networks 1315
Woo-Hyun Ko, *Texas A&M University*; Ushasi Ghosh, *University of California San Diego*; Ujwal Dinesha,
Texas A&M University; Raini Wu, *University of California San Diego*; Srinivas Shakkottai, *Texas A&M University*;
Dinesh Bharadia, *University of California San Diego*

ADR-X: ANN-Assisted Wireless Link Rate Adaptation for Compute-Constrained Embedded Gaming Devices .. 1331
Hao Yin, *University of Washington*; Murali Ramanujam, *Princeton University*; Joe Schaefer, Stan Adermann, Srihari Narlanka,
and Perry Lea, *Microsoft*; Ravi Netravali, *Princeton University*; Krishna Chintalapudi, *Microsoft Research*

RFID+: Spatially Controllable Identification of UHF RFIDs via Controlled Magnetic Fields 1351
Donghui Dai, *The Hong Kong Polytechnic University*; Zhenlin An, *The Hong Kong Polytechnic University and Princeton
University*; Zheng Gong, *The Hong Kong Polytechnic University*; Qingrui Pan, *The Hong Kong Polytechnic University
and The University of Edinburgh*; Lei Yang, *Shenzhen Research Institute, The Hong Kong Polytechnic University*

SMUFF: Towards Line Rate Wi-Fi Direct Transport with Orchestrated On-device Buffer Management 1369
Chengke Wang, *Peking University*; Hao Wang, *Shenzhen Kaihong Digital Industry Development Co., Ltd.*; Yuhan Zhou
and Yunzhe Ni, *Peking University*; Feng Qian, *University of Southern California*; Chenren Xu, *Peking University,
Zhongguancun Laboratory, and Key Laboratory of High Confidence Software Technologies, Ministry of Education (PKU)*

Thursday, April 18

ML Scheduling

Vulcan: Automatic Query Planning for Live ML Analytics 1385
Yiwen Zhang and Xumiao Zhang, *University of Michigan*; Ganesh Ananthanarayanan, *Microsoft*; Anand Iyer,
Georgia Institute of Technology; Yuanhao Shu, *Zhejiang University*; Victor Bahl, *Microsoft Corporation*;
Z. Morley Mao, *University of Michigan and Google*; Mosharaf Chowdhury, *University of Michigan*

CASSINI: Network-Aware Job Scheduling in Machine Learning Clusters 1403
Sudarsanan Rajasekaran and Manya Ghobadi, *Massachusetts Institute of Technology*; Aditya Akella, *UT Austin*

Towards Domain-Specific Network Transport for Distributed DNN Training1421
Hao Wang and Han Tian, *iSING Lab, Hong Kong University of Science and Technology*; Jingrong Chen, *Duke University*;
Xinchen Wan, Jiacheng Xia, and Gaoxiong Zeng, *iSING Lab, Hong Kong University of Science and Technology*;
Wei Bai, *Microsoft*; Junchen Jiang, *University of Chicago*; Yong Wang and Kai Chen, *iSING Lab, Hong Kong University
of Science and Technology*

Swing: Short-cutting Rings for Higher Bandwidth Allreduce 1445
Daniele De Sensi, *Sapienza University of Rome*; Tommaso Bonato, *ETH Zurich*; David Saam, *RWTH Aachen University*;
Torsten Hoefler, *ETH Zurich*

Cloud Scheduling

LitePred: Transferable and Scalable Latency Prediction for Hardware-Aware Neural Architecture Search 1463
Chengquan Feng, *University of Science and Technology of China*; Li Lyna Zhang, *Microsoft Research*; Yuanchi Liu,
University of Science and Technology of China; Jiahang Xu and Chengruidong Zhang, *Microsoft Research*; Zhiyuan Wang,
University of Science and Technology of China; Ting Cao and Mao Yang, *Microsoft Research*; Haisheng Tan, *University of
Science and Technology of China*

Harmonic: Hardware-assisted RDMA Performance Isolation for Public Clouds1479
Jiaqi Lou, *University of Illinois Urbana-Champaign*; Xinhao Kong, *Duke University*; Jinghan Huang, *University of
Illinois Urbana-Champaign*; Wei Bai, *Microsoft*; Nam Sung Kim, *University of Illinois Urbana-Champaign*;
Danyang Zhuo, *Duke University*

LDB: An Efficient Latency Profiling Tool for Multithreaded Applications 1497
Inho Cho, *MIT CSAIL*; Seo Jin Park, *University of Southern California*; Ahmed Saeed, *Georgia Tech*;
Mohammad Alizadeh and Adam Belay, *MIT CSAIL*

UFO: The Ultimate QoS-Aware Core Management for Virtualized and Oversubscribed Public Clouds1511
Yajuan Peng, *Southern University of Science and Technology and Shenzhen Institutes of Advanced Technology,
Chinese Academy of Science*; Shuang Chen and Yi Zhao, *Shuhai Lab, Huawei Cloud*; Zhibin Yu, *Shuhai Lab,
Huawei Cloud, and Shenzhen Institutes of Advanced Technology, Chinese Academy of Science*

Programming the Network: Part 2

Automatic Parallelization of Software Network Functions 1531
Francisco Pereira, Fernando M.V. Ramos, and Luis Pedrosa, *INESC-ID, Instituto Superior Técnico, University of Lisbon*

AutoSketch: Automatic Sketch-Oriented Compiler for Query-driven Network Telemetry 1551
Haifeng Sun and Qun Huang, *National Key Laboratory for Multimedia Information Processing, School of Computer Science,
Peking University*; Jinbo Sun, *Institute of Computing Technology, Chinese Academy of Sciences*; Wei Wang, *Northeastern
University, China*; Jiaheng Li, *National Key Laboratory for Multimedia Information Processing, School of Computer Science,
Peking University*; Fuliang Li, *Northeastern University, China*; Yungang Bao, *Institute of Computing Technology, Chinese
Academy of Sciences*; Xin Yao and Gong Zhang, *Huawei Theory Department*

Leo: Online ML-based Traffic Classification at Multi-Terabit Line Rate 1573
Syed Usman Jafri, Sanjay Rao, Vishal Shrivastav, and Mohit Tawarmalani, *Purdue University*

Sequence Abstractions for Flexible, Line-Rate Network Monitoring 1593
Andrew Johnson, *Princeton University*; Ryan Beckett, *Microsoft Research*; Xiaoqi Chen, *Princeton University*;
Ratul Mahajan, *University of Washington*; David Walker, *Princeton University*

OctoSketch: Enabling Real-Time, Continuous Network Monitoring over Multiple Cores1621
Yinda Zhang, *University of Pennsylvania*; Peiqing Chen and Zaoxing Liu, *University of Maryland*

Wireless Sensing

NR-Surface: NextG-ready μ W-reconfigurable mmWave Metasurface1641
Minseok Kim, Namjo Ahn, and Song Min Kim, *KAIST*

**Cyclops: A Nanomaterial-based, Battery-Free Intraocular Pressure (IOP) Monitoring System
inside Contact Lens** 1659
Liyao Li, *University at Buffalo SUNY and Northwest University*; Bozhao Shang and Yun Wu, *Northwest University
and Shaanxi International Joint Research Centre for the Battery-Free Internet of Things*; Jie Xiong, *University of
Massachusetts Amherst and Microsoft Research Asia*; Xiaojiang Chen, *Northwest University and Shaanxi International
Joint Research Centre for the Battery-Free Internet of Things*; Yaxiong Xie, *University at Buffalo SUNY*

Habitus: Boosting Mobile Immersive Content Delivery through Full-body Pose Tracking and Multipath Networking. 1677
Anlan Zhang, *University of Southern California*; Chendong Wang, *University of Wisconsin – Madison*; Yuming Hu, *University of Minnesota – Twin Cities*; Ahmad Hassan and Zejun Zhang, *University of Southern California*; Bo Han, *George Mason University*; Feng Qian, *University of Southern California*; Shichang Xu, *Google*

BFMSense: WiFi Sensing Using Beamforming Feedback Matrix. 1697
Enze Yi and Dan Wu, *Peking University*; Jie Xiong, *University of Massachusetts Amherst*; Fusang Zhang, *Institute of Software, Chinese Academy of Sciences and University of Chinese Academy of Sciences*; Kai Niu, *Beijing Xiaomi Mobile Software Company Ltd.*; Wenwei Li, *Peking University*; Daqing Zhang, *Peking University and Institut Polytechnique de Paris*

mmComb: High-speed mmWave Commodity WiFi Backscatter.1713
Yoon Chae and Zhenzhe Lin, *George Mason University*; Kang Min Bae and Song Min Kim, *Korea Advanced Institute of Science and Technology (KAIST)*; Parth Pathak, *George Mason University*

Security

Where The Wild Things Are: Brute-Force SSH Attacks In The Wild And How To Stop Them.1731
Sachin Kumar Singh and Shreeman Gautam, *University of Utah*; Cameron Cartier, *University of Utah and Black Hills Information Security*; Sameer Patil and Robert Ricci, *University of Utah*

A System to Detect Forged-Origin BGP Hijacks.1751
Thomas Holterbach and Thomas Alfroy, *University of Strasbourg*; Amreesh Phokeer, *Internet Society*; Alberto Dainotti, *Georgia Tech*; Cristel Pelsser, *UCLouvain*

NetVigil: Robust and Low-Cost Anomaly Detection for East-West Data Center Security1771
Kevin Hsieh, *Microsoft*; Mike Wong, *Princeton University and Microsoft*; Santiago Segarra, *Microsoft and Rice University*; Sathiya Kumaran Mani, Trevor Eberl, and Anatoliy Panasyuk, *Microsoft*; Ravi Netravali, *Princeton University*; Ranveer Chandra and Srikanth Kandula, *Microsoft*

TANGO: Secure Collaborative Route Control across the Public Internet1791
Henry Birge-Lee, Sophia Yoo, Benjamin Herber, Jennifer Rexford, and Maria Apostolaki, *Princeton University*

Sidekick: In-Network Assistance for Secure End-to-End Transport Protocols. 1813
Gina Yuan, Matthew Sotoudeh, and David K. Zhang, *Stanford University*; Michael Welzl, *University of Oslo*; David Mazières and Keith Winstein, *Stanford University*

Mobile Things

VILAM: Infrastructure-assisted 3D Visual Localization and Mapping for Autonomous Driving 1831
Jiahe Cui, *Beihang University, The Chinese University of Hong Kong, and Tianmushan Laboratory*; Shuyao Shi and Yuze He, *The Chinese University of Hong Kong*; Jianwei Niu, *Beihang University*; Guoliang Xing, *The Chinese University of Hong Kong*; Zhenchao Ouyang, *Tianmushan Laboratory and International Innovation Institute of Beihang University*

Catch Me If You Can: Laser Tethering with Highly Mobile Targets 1847
Charles J. Carver, Hadleigh Schwartz, and Qijia Shao, *Columbia University*; Nicholas Shade, Joseph Lazzaro, Xiaoxin Wang, Jifeng Liu, and Eric Fossum, *Dartmouth College*; Xia Zhou, *Columbia University*

MobileConfig: Remote Configuration Management for Mobile Apps at Hyperscale 1867
Matt Guo, *Meta Platforms*; Soteris Demetriou, *Imperial College London*; Joey Yang, Michael Leighton, Diedi Hu, Tong Bao, Amit Adhikari, Thawan Kooburat, Annie Kim, and Chunqiang Tang, *Meta Platforms*

Passengers' Safety Matters: Experiences of Deploying a Large-Scale Indoor Delivery Monitoring System 1883
Xiubin Fan, *City University of Hong Kong*; Zhongming Lin, *The Hong Kong University of Science and Technology*; Yuming Hu, *University of Minnesota - Twin Cities*; Tianrui Jiang, *The Hong Kong University of Science and Technology*; Feng Qian, *University of Southern California*; Zhimeng Yin, *City University of Hong Kong*; S.-H. Gary Chan, *The Hong Kong University of Science and Technology*; Dapeng Wu, *City University of Hong Kong*

AUGUR: Practical Mobile Multipath Transport Service for Low Tail Latency in Real-Time Streaming 1901
Yuhan Zhou, *School of Computer Science, Peking University and Tencent Inc.*; Tingfeng Wang, *Tencent Inc.*; Liying Wang, *School of Computer Science, Peking University*; Nian Wen, Rui Han, Jing Wang, Chenglei Wu, Jiafeng Chen, and Longwei Jiang, *Tencent Inc.*; Shibo Wang, *Xi'an Jiaotong University and Tencent Inc.*; Honghao Liu, *Tencent Inc.*; Chenren Xu, *School of Computer Science, Peking University and Zhongguancun Laboratory and Key Laboratory of High Confidence Software Technologies, Ministry of Education (PKU)*

Cloud Systems

Zombie: Middleboxes that Don't Snoop1917
Collin Zhang, *Cornell*; Zachary DeStefano, Arasu Arun, and Joseph Bonneau, *NYU*; Paul Grubbs, *University of Michigan*;
Michael Walfish, *NYU*

Solving Max-Min Fair Resource Allocations Quickly on Large Graphs 1937
Pooria Namyar, *Microsoft and University of Southern California*; Behnaz Arzani and Srikanth Kandula, *Microsoft*;
Santiago Segarra, *Microsoft and Rice University*; Daniel Crankshaw and Umesh Krishnaswamy, *Microsoft*;
Ramesh Govindan, *University of Southern California*; Himanshu Raj, *Microsoft*

Cloud-LoRa: Enabling Cloud Radio Access LoRa Networks Using Reinforcement Learning Based Bandwidth-Adaptive Compression 1959
Muhammad Osama Shahid, Daniel Koch, Jayaram Raghuram, and Bhuvana Krishnaswamy, *University of Wisconsin-Madison*;
Krishna Chintalapudi, *Microsoft Research*; Suman Banerjee, *University of Wisconsin-Madison*

Cloudy with a Chance of Cyberattacks: Dangling Resources Abuse on Cloud Platforms 1977
Jens Frieß, *National Research Center for Applied Cybersecurity ATHENE and Technische Universität Darmstadt*;
Tobias Gattermayer, *National Research Center for Applied Cybersecurity ATHENE and Fraunhofer Institute for Secure Information Technology SIT*; Nethanel Gelernter, *IONIX*; Haya Schulmann, *National Research Center for Applied Cybersecurity ATHENE and Goethe-Universität Frankfurt*; Michael Waidner, *National Research Center for Applied Cybersecurity ATHENE, Technische Universität Darmstadt, and Fraunhofer Institute for Secure Information Technology SIT*

Modeling Networks

CAPA: An Architecture For Operating Cluster Networks With High Availability 1995
Bingzhe Liu, *UIUC*; Colin Scott, Mukarram Tariq, Andrew Ferguson, Phillipa Gill, Richard Alimi, Omid Alipourfard, Deepak Arulkannan, Virginia Jean Beauregard, and Patrick Conner, *Google*; P. Brighten Godfrey, *UIUC*; Xander Lin, Joon Ong, Mayur Patel, Amr Sabaa, Arjun Singh, Alex Smirnov, Manish Verma, Prerepa V Viswanadham, and Amin Vahdat, *Google*

NetAssistant: Dialogue Based Network Diagnosis in Data Center Networks 2011
Haopei Wang, Anubhavnidhi Abhashkumar, Changyu Lin, Tianrong Zhang, Xiaoming Gu, Ning Ma, Chang Wu, Songlin Liu, Wei Zhou, Yongbin Dong, Weirong Jiang, and Yi Wang, *ByteDance Inc*

Klonet: an Easy-to-Use and Scalable Platform for Computer Networks Education 2025
Tie Ma, Long Luo, and Hongfang Yu, *University of Electronic Science and Technology of China*; Xi Chen, *Southwest Minzu University*; Jingzhao Xie, Chongxi Ma, Yunhan Xie, Gang Sun, and Tianxi Wei, *University of Electronic Science and Technology of China*; Li Chen, *Zhongguancun Laboratory*; Yanwei Xu and Nicholas Zhang, *Theory Lab, Central Research Institute, 2012 Labs, Huawei Technologies Co., Ltd.*

EXCHAIN: Exception Dependency Analysis for Root Cause Diagnosis 2047
Ao Li, *Carnegie Mellon University*; Shan Lu, *Microsoft Research and University of Chicago*; Suman Nath, *Microsoft Research*;
Rohan Padhye and Vyas Sekar, *Carnegie Mellon University*