

# **20th USENIX Symposium on Networked Systems Design and Implementation (NSDI'23)**

Boston, Massachusetts, USA  
17-19 April 2023

Volume 1 of 3

ISBN: 978-1-7138-7231-3

**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© (2023) by USENIX Association  
All rights reserved.

Printed with permission by Curran Associates, Inc. (2023)

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](mailto:curran@proceedings.com)  
Web: [www.proceedings.com](http://www.proceedings.com)

# 20th USENIX Symposium on Networked Systems Design and Implementation (NSDI '23)

April 17–19, 2023  
Boston, MA, USA

## Monday, April 17

### RDMA

#### **SRNIC: A Scalable Architecture for RDMA NICs** . . . . . 1

Zilong Wang, *Hong Kong University of Science and Technology*; Layong Luo and Qingsong Ning, *ByteDance*; Chaoliang Zeng, Wenxue Li, and Xinchun Wan, *Hong Kong University of Science and Technology*; Peng Xie, Tao Feng, Ke Cheng, Xiongfei Geng, Tianhao Wang, Weicheng Ling, Kejia Huo, Pingbo An, Kui Ji, Shideng Zhang, Bin Xu, Ruiqing Feng, and Tao Ding, *ByteDance*; Kai Chen, *Hong Kong University of Science and Technology*; Chuanxiong Guo

#### **Hostping: Diagnosing Intra-host Network Bottlenecks in RDMA Servers** . . . . . 15

Kefei Liu, *BUPT*; Zhuo Jiang, *ByteDance Inc.*; Jiao Zhang, *BUPT and Purple Mountain Laboratories*; Haoran Wei, *BUPT and ByteDance Inc.*; Xiaolong Zhong, *BUPT*; Lizhuang Tan, *ByteDance Inc.*; Tian Pan and Tao Huang, *BUPT and Purple Mountain Laboratories*

#### **Understanding RDMA Microarchitecture Resources for Performance Isolation** . . . . . 31

Xinhao Kong and Jingrong Chen, *Duke University*; Wei Bai, *Microsoft*; Yechen Xu, *Shanghai Jiao Tong University*; Mahmoud Elhaddad, Shachar Raindel, and Jitendra Padhye, *Microsoft*; Alvin R. Lebeck and Danyang Zhuo, *Duke University*

#### **Empowering Azure Storage with RDMA** . . . . . 49

Wei Bai, Shanm Sainul Abdeen, Ankit Agrawal, Krishan Kumar Attre, Paramvir Bahl, Ameya Bhagat, Gowri Bhaskara, Tanya Brokhman, Lei Cao, Ahmad Cheema, Rebecca Chow, Jeff Cohen, Mahmoud Elhaddad, Vivek Ette, Igal Figlin, Daniel Firestone, Mathew George, Ilya German, Lakhmeet Ghai, Eric Green, Albert Greenberg, Manish Gupta, Randy Haagens, Matthew Hendel, Ridwan Howlader, Neetha John, Julia Johnstone, Tom Jolly, Greg Kramer, David Kruse, Ankit Kumar, Erica Lan, Ivan Lee, Avi Levy, Marina Lipshteyn, Xin Liu, Chen Liu, Guohan Lu, Yuemin Lu, Xiakun Lu, Vadim Makhervaks, Ulad Malashanka, David A. Maltz, Ilias Marinos, Rohan Mehta, Sharda Murthi, Anup Namdhari, Aaron Ogus, Jitendra Padhye, Madhav Pandya, Douglas Phillips, Adrian Power, Suraj Puri, Shachar Raindel, Jordan Rhee, Anthony Russo, Maneesh Sah, Ali Sheriff, Chris Sparacino, Ashutosh Srivastava, Weixiang Sun, Nick Swanson, Fuhou Tian, Lukasz Tomczyk, Vamsi Vadlamuri, Alec Wolman, Ying Xie, Joyce Yom, Lihua Yuan, Yanzhao Zhang, and Brian Zill, *Microsoft*

### Learning with GPUs

#### **Transparent GPU Sharing in Container Clouds for Deep Learning Workloads** . . . . . 69

Binyang Wu and Zili Zhang, *Peking University*; Zhihao Bai, *Johns Hopkins University*; Xuanzhe Liu and Xin Jin, *Peking University*

#### **ARK: GPU-driven Code Execution for Distributed Deep Learning** . . . . . 87

Changho Hwang, *KAIST, Microsoft Research*; KyoungSoo Park, *KAIST*; Ran Shu, Xinyuan Qu, Peng Cheng, and Yongqiang Xiong, *Microsoft Research*

#### **BGL: GPU-Efficient GNN Training by Optimizing Graph Data I/O and Preprocessing** . . . . . 103

Tianfeng Liu, *Tsinghua University, Zhongguancun Laboratory, ByteDance*; Yangrui Chen, *The University of Hong Kong, ByteDance*; Dan Li, *Tsinghua University, Zhongguancun Laboratory*; Chuan Wu, *The University of Hong Kong*; Yibo Zhu, Jun He, and Yanghua Peng, *ByteDance*; Hongzheng Chen, *ByteDance, Cornell University*; Hongzhi Chen and Chuanxiong Guo, *ByteDance*

#### **Zeus: Understanding and Optimizing GPU Energy Consumption of DNN Training** . . . . . 119

Jie You, Jae-Won Chung, and Mosharaf Chowdhury, *University of Michigan*

## RPC and Remote Memory

- Remote Procedure Call as a Managed System Service** .....141  
Jingrong Chen, Yongji Wu, and Shihan Lin, *Duke University*; Yechen Xu, *Shanghai Jiao Tong University*; Xinhao Kong, *Duke University*; Thomas Anderson, *University of Washington*; Matthew Lentz, Xiaowei Yang, and Danyang Zhuo, *Duke University*
- Canvas: Isolated and Adaptive Swapping for Multi-Applications on Remote Memory** .....161  
Chenxi Wang, Yifan Qiao, Haoran Ma, and Shi Liu, *UCLA*; Yiying Zhang, *UCSD*; Wenguang Chen, *Tsinghua University*; Ravi Netravali, *Princeton University*; Miryung Kim and Guoqing Harry Xu, *UCLA*
- Hermit: Low-Latency, High-Throughput, and Transparent Remote Memory via Feedback-Directed Asynchrony** ...181  
Yifan Qiao and Chenxi Wang, *UCLA*; Zhenyuan Ruan and Adam Belay, *MIT CSAIL*; Qingda Lu, *Alibaba Group*; Yiying Zhang, *UCSD*; Miryung Kim and Guoqing Harry Xu, *UCLA*
- NetRPC: Enabling In-Network Computation in Remote Procedure Calls** ..... 199  
Bohan Zhao, *Tsinghua University*; Wenfei Wu, *Peking University*; Wei Xu, *Tsinghua University*

## Congestion Control

- Bolt: Sub-RTT Congestion Control for Ultra-Low Latency** ..... 219  
Serhat Arslan, *Stanford University*; Yuliang Li, Gautam Kumar, and Nandita Dukkhipati, *Google LLC*
- Understanding the impact of host networking elements on traffic bursts** ..... 237  
Erfan Sharafzadeh and Sepehr Abdous, *Johns Hopkins University*; Soudeh Ghorbani, *Johns Hopkins University and Meta*
- Poseidon: Efficient, Robust, and Practical Datacenter CC via Deployable INT** ..... 255  
Weitao Wang, *Google LLC and Rice University*; Masoud Moshref, Yuliang Li, and Gautam Kumar, *Google LLC*; T. S. Eugene Ng, *Rice University*; Neal Cardwell and Nandita Dukkhipati, *Google LLC*
- Rearchitecting the TCP Stack for I/O-Offloaded Content Delivery** ..... 275  
Taehyun Kim and Deondre Martin Ng, *KAIST*; Junzhi Gong, *Harvard University*; Youngjin Kwon, *KAIST*; Minlan Yu, *Harvard University*; Kyoungsoo Park, *KAIST*

## Distributed Systems

- Hydra: Serialization-Free Network Ordering for Strongly Consistent Distributed Applications** ..... 293  
Inho Choi, *National University of Singapore*; Ellis Michael, *University of Washington*; Yunfan Li, *National University of Singapore*; Dan R. K. Ports, *Microsoft Research*; Jialin Li, *National University of Singapore*
- The Benefit of Hindsight: Tracing Edge-Cases in Distributed Systems** ..... 321  
Lei Zhang, *Emory University and Princeton University*; Zhiqiang Xie and Vaastav Anand, *Max Planck Institute for Software Systems*; Ymir Vigfusson, *Emory University*; Jonathan Mace, *Max Planck Institute for Software Systems*
- DiSh: Dynamic Shell-Script Distribution** ..... 341  
Tammam Mustafa, *MIT*; Konstantinos Kallas, *University of Pennsylvania*; Pratyush Das, *Purdue University*; Nikos Vasilakis, *Brown University*
- Waverunner: An Elegant Approach to Hardware Acceleration of State Machine Replication** ..... 357  
Mohammadreza Alimadadi and Hieu Mai, *Stony Brook University*; Shengsun Cho, *Microsoft*; Michael Ferdman, Peter Milder, and Shuai Mu, *Stony Brook University*

## Wireless

- LeakyScatter: A Frequency-Agile Directional Backscatter Network Above 100 GHz** ..... 375  
Atsute Kludze and Yasaman Ghasempour, *Princeton University*
- RF-Bouncer: A Programmable Dual-band Metasurface for Sub-6 Wireless Networks** ..... 389  
Xinyi Li, Chao Feng, Xiaojing Wang, and Yangfan Zhang, *Northwest University*; Yaxiong Xie, *University at Buffalo SUNY*; Xiaojiang Chen, *Northwest University*
- Scalable Distributed Massive MIMO Baseband Processing** ..... 405  
Junzhi Gong, *Harvard University*; Anuj Kalia, *Microsoft*; Minlan Yu, *Harvard University*

**DChannel: Accelerating Mobile Applications With Parallel High-bandwidth and Low-latency Channels** . . . . . 419  
William Sentosa, *University of Illinois Urbana-Champaign*; Balakrishnan Chandrasekaran, *Vrije Universiteit Amsterdam*;  
P. Brighten Godfrey, *University of Illinois Urbana-Champaign and VMware*; Haitham Hassanieh, *EPFL*; Bruce Maggs,  
*Duke University and Emerald Innovations*

## Cloud

**SkyPilot: An Intercloud Broker for Sky Computing** . . . . . 437  
Zongheng Yang, Zhanghao Wu, Michael Luo, Wei-Lin Chiang, Romil Bhardwaj, Woosuk Kwon, Siyuan Zhuang,  
Frank Sifei Luan, and Gautam Mittal, *UC Berkeley*; Scott Shenker, *UC Berkeley and ICSI*; Ion Stoica, *UC Berkeley*

**Unlocking unallocated cloud capacity for long, uninterruptible workloads** . . . . . 457  
Anup Agarwal, *Carnegie Mellon University*; Shadi Noghahi, *Microsoft Research*; Íñigo Goiri, *Azure Systems Research*;  
Srinivasan Seshan, *Carnegie Mellon University*; Anirudh Badam, *Microsoft Research*

**Invisinets: Removing Networking from Cloud Networks** . . . . . 479  
Sarah McClure and Zeke Medley, *UC Berkeley*; Deepak Bansal and Karthick Jayaraman, *Microsoft*; Ashok Narayanan,  
*Google*; Jitendra Padhye, *Microsoft*; Sylvia Ratnasamy, *UC Berkeley and Google*; Anees Shaikh, *Google*; Rishabh Tewari,  
*Microsoft*

**Bamboo: Making Preemptible Instances Resilient for Affordable Training of Large DNNs** . . . . . 497  
John Thorpe, Pengzhan Zhao, Jonathan Eyolfson, and Yifan Qiao, *UCLA*; Zhihao Jia, *CMU*; Minjia Zhang, *Microsoft  
Research*; Ravi Netravali, *Princeton University*; Guoqing Harry Xu, *UCLA*

## Internet-Scale Networks

**ONEWAN is better than two: Unifying a split WAN architecture** . . . . . 515  
Umesh Krishnaswamy, *Microsoft*; Rachee Singh, *Microsoft and Cornell University*; Paul Mattes, Paul-Andre  
C Bissonnette, Nikolaj Bjørner, Zahira Nasrin, Sonal Kothari, Prabhakar Reddy, John Abeln, Srikanth Kandula,  
Himanshu Raj, Luis Irun-Briz, Jamie Gaudette, and Erica Lan, *Microsoft*

**RHINE: Robust and High-performance Internet Naming with E2E Authenticity** . . . . . 531  
Huayi Duan, Rubén Fischer, Jie Lou, Si Liu, David Basin, and Adrian Perrig, *ETH Zürich*

**Enabling Users to Control their Internet** . . . . . 555  
Ammar Tahir and Radhika Mittal, *University of Illinois at Urbana-Champaign*

**xBGP: Faster Innovation in Routing Protocols** . . . . . 575  
Thomas Wirtgen, Tom Rousseaux, Quentin De Coninck, and Nicolas Rybowski, *ICTEAM, UCLouvain*; Randy Bush,  
*Internet Initiative Japan & Arrcus, Inc*; Laurent Vanbever, *NSG, ETH Zürich*; Axel Legay and Olivier Bonaventure,  
*ICTEAM, UCLouvain*

## Tuesday, April 18

### Synthesis and Formal Methods

**TACCL: Guiding Collective Algorithm Synthesis using Communication Sketches** . . . . . 593  
Aashaka Shah, *University of Texas at Austin*; Vijay Chidambaram, *University of Texas at Austin and VMware Research*;  
Meghan Cowan, Saeed Maleki, Madan Musuvathi, Todd Mytkowicz, Jacob Nelson, and Olli Saarikivi, *Microsoft Research*;  
Rachee Singh, *Microsoft and Cornell University*

**Synthesizing Runtime Programmable Switch Updates** . . . . . 613  
Yiming Qiu, *Rice University*; Ryan Beckett, *Microsoft*; Ang Chen, *Rice University*

**Practical Intent-driven Routing Configuration Synthesis** . . . . . 629  
Sivaramakrishnan Ramanathan, Ying Zhang, Mohab Gawish, Yogesh Mundada, Zhaodong Wang, Sangki Yun,  
Eric Lippert, and Walid Taha, *Meta*; Minlan Yu, *Harvard University*; Jelena Mirkovic, *University of Southern California  
Information Sciences Institute*

**Formal Methods for Network Performance Analysis** . . . . . 645  
Mina Tahmasbi Arashloo, *University of Waterloo*; Ryan Beckett, *Microsoft Research*; Rachit Agarwal, *Cornell University*

## Data Centers

- Flattened Clos: Designing High-performance Deadlock-free Expander Data Center Networks Using Graph Contraction** ..... 663  
Shizhen Zhao, Qizhou Zhang, Peirui Cao, Xiao Zhang, and Xinbing Wang, *Shanghai Jiao Tong University*;  
Chenghu Zhou, *Shanghai Jiao Tong University and Chinese Academy of Sciences*
- Scalable Tail Latency Estimation for Data Center Networks**..... 685  
Kevin Zhao, *University of Washington*; Prateesh Goyal, *Microsoft Research*; Mohammad Alizadeh, *MIT CSAIL*;  
Thomas E. Anderson, *University of Washington*
- Shockwave: Fair and Efficient Cluster Scheduling for Dynamic Adaptation in Machine Learning** ..... 703  
Pengfei Zheng and Rui Pan, *University of Wisconsin-Madison*; Tarannum Khan, *The University of Texas at Austin*;  
Shivaram Venkataraman, *University of Wisconsin-Madison*; Aditya Akella, *The University of Texas at Austin*
- Protego: Overload Control for Applications with Unpredictable Lock Contention**..... 725  
Inho Cho, *MIT CSAIL*; Ahmed Saeed, *Georgia Tech*; Seo Jin Park, Mohammad Alizadeh, and Adam Belay, *MIT CSAIL*

## Systems for Learning

- TopoOpt: Co-optimizing Network Topology and Parallelization Strategy for Distributed Training Jobs**..... 739  
Weiyang Wang, Moein Khazraee, Zhizhen Zhong, and Manya Ghobadi, *Massachusetts Institute of Technology*;  
Zhihao Jia, *Meta and CMU*; Dheevatsa Mudigere and Ying Zhang, *Meta*; Anthony Kewitsch, *Telescent*
- ModelKeeper: Accelerating DNN Training via Automated Training Warmup** ..... 769  
Fan Lai, Yinwei Dai, Harsha V. Madhyastha, and Mosharaf Chowdhury, *University of Michigan*
- SHEPHERD: Serving DNNs in the Wild** ..... 787  
Hong Zhang, *University of Waterloo*; Yupeng Tang and Anurag Khandelwal, *Yale University*; Ion Stoica, *UC Berkeley*
- Better Together: Jointly Optimizing ML Collective Scheduling and Execution Planning using SYNDICATE** ..... 809  
Kshiteej Mahajan, *University of Wisconsin - Madison*; Ching-Hsiang Chu and Srinivas Sridharan, *Facebook*;  
Aditya Akella, *UT Austin*

## Privacy and Security

- Addax: A fast, private, and accountable ad exchange infrastructure** ..... 825  
Ke Zhong, Yiping Ma, and Yifeng Mao, *University of Pennsylvania*; Sebastian Angel, *University of Pennsylvania & Microsoft Research*
- SPEEDEX: A Scalable, Parallelizable, and Economically Efficient Decentralized EXchange** ..... 849  
Geoffrey Ramseyer, Ashish Goel, and David Mazières, *Stanford University*
- Boomerang: Metadata-Private Messaging under Hardware Trust** ..... 877  
Peipei Jiang, *Wuhan University and City University of Hong Kong*; Qian Wang and Jianhao Cheng, *Wuhan University*;  
Cong Wang, *City University of Hong Kong*; Lei Xu, *Nanjing University of Science and Technology*; Xinyu Wang, *Tencent Inc.*; Yihao Wu and Xiaoyuan Li, *Wuhan University*; Kui Ren, *Zhejiang University*
- Hamilton: A High-Performance Transaction Processor for Central Bank Digital Currencies** ..... 901  
James Lovejoy, *Federal Reserve Bank of Boston*; Madars Virza and Cory Fields, *MIT Media Lab*; Kevin Karwaski and Anders Brownworth, *Federal Reserve Bank of Boston*; Neha Narula, *MIT Media Lab*

## Video

- RECL: Responsive Resource-Efficient Continuous Learning for Video Analytics** .....917  
Mehrdad Khani, *MIT CSAIL and Microsoft*; Ganesh Ananthanarayanan and Kevin Hsieh, *Microsoft*; Junchen Jiang, *University of Chicago*; Ravi Netravali, *Princeton University*; Yuanchao Shu, *Zhejiang University*; Mohammad Alizadeh, *MIT CSAIL*; Victor Bahl, *Microsoft*
- Boggart: Towards General-Purpose Acceleration of Retrospective Video Analytics** ..... 933  
Neil Agarwal and Ravi Netravali, *Princeton University*



**Tambur: Efficient loss recovery for videoconferencing via streaming codes** ..... 953  
Michael Rudow, *Carnegie Mellon University*; Francis Y. Yan, *Microsoft Research*; Abhishek Kumar, *Carnegie Mellon University*; Ganesh Ananthanarayanan and Martin Ellis, *Microsoft*; K.V. Rashmi, *Carnegie Mellon University*

**Gemel: Model Merging for Memory-Efficient, Real-Time Video Analytics at the Edge** ..... 973  
Arthi Padmanabhan, *UCLA*; Neil Agarwal, *Princeton University*; Anand Iyer and Ganesh Ananthanarayanan, *Microsoft Research*; Yuanchao Shu, *Zhejiang University*; Nikolaos Karianakis, *Microsoft Research*; Guoqing Harry Xu, *UCLA*; Ravi Netravali, *Princeton University*

## Data

**Fast, Approximate Vector Queries on Very Large Unstructured Datasets** ..... 995  
Zili Zhang and Chao Jin, *Peking University*; Linpeng Tang, *Moqi*; Xuanzhe Liu and Xin Jin, *Peking University*

**Arya: Arbitrary Graph Pattern Mining with Decomposition-based Sampling** .....1013  
Zeyang Zhu, *Boston University*; Kan Wu, *University of Wisconsin-Madison*; Zaoxing Liu, *Boston University*

**SECRECY: Secure collaborative analytics in untrusted clouds** .....1031  
John Liagouris, Vasiliki Kalavri, Muhammad Faisal, and Mayank Varia, *Boston University*

**FLASH: Towards a High-performance Hardware Acceleration Architecture for Cross-silo Federated Learning** 1057  
Junxue Zhang and Xiaodian Cheng, *iSINGLab at Hong Kong University of Science and Technology and Cluster*; Wei Wang, *Cluster*; Liu Yang, *iSINGLab at Hong Kong University of Science and Technology and Cluster*; Jinbin Hu and Kai Chen, *iSINGLab at Hong Kong University of Science and Technology*

## Making Systems Learn

**On Modular Learning of Distributed Systems for Predicting End-to-End Latency** ..... 1081  
Chieh-Jan Mike Liang, *Microsoft Research*; Zilin Fang, *Carnegie Mellon University*; Yuqing Xie, *Tsinghua University*; Fan Yang, *Microsoft Research*; Zhao Lucis Li, *University of Science and Technology of China*; Li Lina Zhang, Mao Yang, and Lidong Zhou, *Microsoft Research*

**SelfTune: Tuning Cluster Managers** ..... 1097  
Ajaykrishna Karthikeyan and Nagarajan Natarajan, *Microsoft Research*; Gagan Somashekar, *Stony Brook University*; Lei Zhao, *Microsoft*; Ranjita Bhagwan, *Microsoft Research*; Rodrigo Fonseca, Tatiana Racheva, and Yogesh Bansal, *Microsoft*

**CausalSim: A Causal Framework for Unbiased Trace-Driven Simulation** .....1115  
Abdullah Alomar, Pouya Hamadani, Arash Nasr-Esfahany, Anish Agarwal, Mohammad Alizadeh, and Devavrat Shah, *MIT*

**HALP: Heuristic Aided Learned Preference Eviction Policy for YouTube Content Delivery Network** .....1149  
Zhenyu Song, *Princeton University*; Kevin Chen, Nikhil Sarda, Deniz Altınbüken, Eugene Brevdo, Jimmy Coleman, Xiao Ju, Pawel Jurczyk, Richard Schooler, and Ramki Gummadi, *Google*

## IoT Networks

**OpenLoRa: Validating LoRa Implementations through an Extensible and Open-sourced Framework** .....1165  
Manan Mishra, Daniel Koch, Muhammad Osama Shahid, and Bhuvana Krishnaswamy, *University of Wisconsin-Madison*; Krishna Chintalapudi, *Microsoft Research*; Suman Banerjee, *University of Wisconsin-Madison*

**VECARE: Statistical Acoustic Sensing for Automotive In-Cabin Monitoring** ..... 1185  
Yi Zhang, *The University of Hong Kong and Tsinghua University*; Weiying Hou, *The University of Hong Kong*; Zheng Yang, *Tsinghua University*; Chenshu Wu, *The University of Hong Kong*

**SlimWiFi: Ultra-Low-Power IoT Radio Architecture Enabled by Asymmetric Communication** ..... 1201  
Renjie Zhao, *University of California San Diego*; Kejia Wang, *Baylor University*; Kai Zheng and Xinyu Zhang, *University of California San Diego*; Vincent Leung, *Baylor University*

**SLNet: A Spectrogram Learning Neural Network for Deep Wireless Sensing** ..... 1221  
Zheng Yang and Yi Zhang, *Tsinghua University*; Kun Qian, *University of California San Diego*; Chenshu Wu, *The University of Hong Kong*

## Wednesday, April 19

### Programming the Network

- A High-Speed Stateful Packet Processing Approach for Tbps Programmable Switches** ..... 1237  
Mariano Scazzariello and Tommaso Caiazz, *KTH Royal Institute of Technology and Roma Tre University*;  
Hamid Ghasemirahni, *KTH Royal Institute of Technology*; Tom Barbette, *UCLouvain*; Dejan Kostić and  
Marco Chiesa, *KTH Royal Institute of Technology*
- ExoPlane: An Operating System for On-Rack Switch Resource Augmentation** ..... 1257  
Daehyeok Kim, *Microsoft and University of Texas at Austin*; Vyas Sekar and Srinivasan Seshan, *Carnegie Mellon University*
- Sketchovsky: Enabling Ensembles of Sketches on Programmable Switches** ..... 1273  
Hun Namkung, *Carnegie Mellon University*; Zaoxing Liu, *Boston University*; Daehyeok Kim, *Microsoft Research*;  
Vyas Sekar and Peter Steenkiste, *Carnegie Mellon University*
- RingLeader: Efficiently Offloading Intra-Server Orchestration to NICs** ..... 1293  
Jiaxin Lin, Adney Cardoza, Tarannum Khan, and Yeonju Ro, *UT Austin*; Brent E. Stephens, *University of Utah*;  
Hassan Wassel, *Google*; Aditya Akella, *UT Austin*

### Alternative Networks

- STARRYNET: Empowering Researchers to Evaluate Futuristic Integrated Space and Terrestrial Networks** ..... 1309  
Zeqi Lai and Hewu Li, *Tsinghua University and Zhongguancun Laboratory*; Yangtao Deng, *Tsinghua University*;  
Qian Wu, Jun Liu, and Yuanjie Li, *Tsinghua University and Zhongguancun Laboratory*; Jihao Li, Lixin Liu, and  
Weisen Liu, *Tsinghua University*; Jianping Wu, *Tsinghua University and Zhongguancun Laboratory*
- POLYCORN: Data-driven Cross-layer Multipath Networking for High-speed Railway through Composable Schedulerlets** ..... 1325  
Yunzhe Ni, *Peking University*; Feng Qian, *University of Minnesota – Twin Cities*; Taide Liu, Yihua Cheng, Zhiyao Ma,  
and Jing Wang, *Peking University*; Zhongfeng Wang, *China Railway Gecent Technology Co., Ltd*; Gang Huang and  
Xuanzhe Liu, *Key Laboratory of High Confidence Software Technologies, Ministry of Education, Peking University*;  
Chenren Xu, *Zhongguancun Laboratory and Key Laboratory of High Confidence Software Technologies, Ministry of  
Education, Peking University*
- Augmenting Augmented Reality with Non-Line-of-Sight Perception** ..... 1341  
Tara Boroushaki, Maisy Lam, and Laura Dodds, *Massachusetts Institute of Technology*; Aline Eid, *Massachusetts  
Institute of Technology and University of Michigan*; Fadel Adib, *Massachusetts Institute of Technology*
- Acoustic Sensing and Communication Using Metasurface** ..... 1359  
Yongzhao Zhang, Yezhou Wang, and Lanqing Yang, *Shanghai Jiao Tong University*; Mei Wang, *UT Austin*; Yi-Chao Chen,  
*Shanghai Jiao Tong University and Microsoft Research Asia*; Lili Qiu, *UT Austin and Microsoft Research Asia*;  
Yihong Liu, *University of Glasgow*; Guangtao Xue and Jiadi Yu, *Shanghai Jiao Tong University*

### Performance

- Skyplane: Optimizing Transfer Cost and Throughput Using Cloud-Aware Overlays** ..... 1375  
Paras Jain, Sam Kumar, Sarah Wooders, Shishir G. Patil, Joseph E. Gonzalez, and Ion Stoica, *University of California,  
Berkeley*
- Electrode: Accelerating Distributed Protocols with eBPF** ..... 1391  
Yang Zhou, *Harvard University*; Zezhou Wang, *Peking University*; Sowmya Dharanipragada, *Cornell University*;  
Minlan Yu, *Harvard University*
- Nu: Achieving Microsecond-Scale Resource Fungibility with Logical Processes** ..... 1409  
Zhenyuan Ruan and Seo Jin Park, *MIT CSAIL*; Marcos K. Aguilera, *VMware Research*; Adam Belay, *MIT CSAIL*;  
Malte Schwarzkopf, *Brown University*
- Enabling High Quality Real-Time Communications with Adaptive Frame-Rate** ..... 1429  
Zili Meng, *Tsinghua University and Tencent Inc.*; Tingfeng Wang, *Tsinghua University, Tencent Inc., and  
Beijing University of Posts and Telecommunications*; Yixin Shen, *Tsinghua University*; Bo Wang and Mingwei Xu,  
*Tsinghua University and Zhongguancun Laboratory*; Rui Han and Honghao Liu, *Tencent Inc.*; Venkat Arun,  
*Massachusetts Institute of Technology*; Hongxin Hu, *University at Buffalo, SUNY*; Xue Wei, *Tencent Inc.*



## Serverless and Network Functions

- LemonNFV: Consolidating Heterogeneous Network Functions at Line Speed** .....1451  
Hao Li and Yihan Dang, *Xi'an Jiaotong University*; Guangda Sun, *Xi'an Jiaotong University and National University of Singapore*; Guyue Liu, *New York University Shanghai*; Danfeng Shan and Peng Zhang, *Xi'an Jiaotong University*
- Disaggregating Stateful Network Functions** ..... 1469  
Deepak Bansal, Gerald DeGrace, Rishabh Tewari, Michal Zygmunt, and James Grantham, *Microsoft*; Silvano Gai, Mario Baldi, Krishna Doddapaneni, Arun Selvarajan, Arunkumar Arumugam, and Balakrishnan Raman, *AMD Pensando*; Avijit Gupta, Sachin Jain, Deven Jagasia, Evan Langlais, Pranjal Srivastava, Rishiraj Hazarika, Neeraj Motwani, Soumya Tiwari, Stewart Grant, Ranveer Chandra, and Srikanth Kandula, *Microsoft*
- Following the Data, Not the Function: Rethinking Function Orchestration in Serverless Computing** ..... 1489  
Minchen Yu, *Hong Kong University of Science and Technology*; Tingjia Cao, *University of Wisconsin-Madison*; Wei Wang, *Hong Kong University of Science and Technology*; Ruichuan Chen, *Nokia Bell Labs*
- Doing More with Less: Orchestrating Serverless Applications without an Orchestrator** ..... 1505  
David H. Liu and Amit Levy, *Princeton University*; Shadi Noghabi and Sebastian Burckhardt, *Microsoft Research*

## Real Networks

- Enhancing Global Network Monitoring with *Magnifier*** ..... 1521  
Tobias Bühler and Romain Jacob, *ETH Zürich*; Ingmar Poesse, *BENOCS*; Laurent Vanbever, *ETH Zürich*
- NetPanel: Traffic Measurement of Exchange Online Service** ..... 1541  
Yu Chen, *Microsoft 365, China*; Liqun Li and Yu Kang, *Microsoft Research, China*; Boyang Zheng, Yehan Wang, More Zhou, Yuchao Dai, and Zhenguo Yang, *Microsoft 365, China*; Brad Rutkowski and Jeff Mealiffe, *Microsoft 365, USA*; Qingwei Lin, *Microsoft Research, China*
- DOTe: Rethinking (Predictive) WAN Traffic Engineering** ..... 1557  
Yarin Perry, *Hebrew University of Jerusalem*; Felipe Vieira Frujeri, *Microsoft Research*; Chaim Hoch, *Hebrew University of Jerusalem*; Srikanth Kandula and Ishai Menache, *Microsoft Research*; Michael Schapira, *Hebrew University of Jerusalem*; Aviv Tamar, *Technion*
- Dashlet: Taming Swipe Uncertainty for Robust Short Video Streaming** ..... 1583  
Zhuqi Li, Yaxiong Xie, Ravi Netravali, and Kyle Jamieson, *Princeton University*

## Cellular

- CellDAM: User-Space, Rootless Detection and Mitigation for 5G Data Plane** .....1601  
Zhaowei Tan, Jinghao Zhao, Boyan Ding, and Songwu Lu, *University of California, Los Angeles*
- LOCA: A Location-Oblivious Cellular Architecture** .....1621  
Zhihong Luo, Silvery Fu, and Natacha Crooks, *UC Berkeley*; Shaddi Hasan, *Virginia Tech*; Christian Maciocco, *Intel*; Sylvia Ratnasamy, *UC Berkeley*; Scott Shenker, *UC Berkeley and ICSI*
- mmWall: A Steerable, Transflective Metamaterial Surface for NextG mmWave Networks** ..... 1647  
Kun Woo Cho, *Princeton University*; Mohammad H. Mazaheri, *UCLA*; Jeremy Gummesson, *University of Massachusetts Amherst*; Omid Abari, *UCLA*; Kyle Jamieson, *Princeton University*
- Building Flexible, Low-Cost Wireless Access Networks With Magma** ..... 1667  
Shaddi Hasan, *Virginia Tech*; Amar Padmanabhan, *Databricks*; Bruce Davie, *Systems Approach*; Jennifer Rexford, *Princeton University*; Ulas Kozat, Hunter Gatewood, Shruti Sanadhya, Nick Yurchenko, Tariq Al-Khasib, Oriol Batalla, Marie Bremner, Andrei Lee, Evgeniy Makeev, Scott Moeller, Alex Rodriguez, Pravin Shelar, Karthik Subraveti, Sudarshan Kandi, Alejandro Xoconostle, and Praveen Kumar Ramakrishnan, *Meta*; Xiaochen Tian, *Independent*; Anoop Tomar, *Meta*

## Testing

**LinkLab 2.0: A Multi-tenant Programmable IoT Testbed for Experimentation with Edge-Cloud Integration ... 1683**  
Wei Dong, Borui Li, Haoyu Li, Hao Wu, Kaijie Gong, Wenzhao Zhang, and Yi Gao, *Zhejiang University*

**Push-Button Reliability Testing for Cloud-Backed Applications with Rainmaker .....1701**  
Yinfang Chen and Xudong Sun, *University of Illinois at Urbana-Champaign*; Suman Nath, *Microsoft Research*;  
Ze Yang and Tianyin Xu, *University of Illinois at Urbana-Champaign*

**Test Coverage for Network Configurations .....1717**  
Xieyang Xu and Weixin Deng, *University of Washington*; Ryan Beckett, *Microsoft*; Ratul Mahajan, *University of Washington*; David Walker, *Princeton University*

**Norma: Towards Practical Network Load Testing .....1733**  
Yanqing Chen, *State Key Laboratory for Novel Software Technology, Nanjing University and Alibaba Group*;  
Bingchuan Tian, *Alibaba Group*; Chen Tian, *State Key Laboratory for Novel Software Technology, Nanjing University*;  
Li Dai, Yu Zhou, Mengjing Ma, and Ming Tang, *Alibaba Group*; Hao Zheng, Zhewen Yang, and Guihai Chen, *State Key Laboratory for Novel Software Technology, Nanjing University*; Dennis Cai and Ennan Zhai, *Alibaba Group*

## Physical Layer

**$\mu$ Mote: Enabling Passive Chirp De-spreading and  $\mu$ W-level Long-Range Downlink for Backscatter Devices ....1751**  
Yihang Song and Li Lu, *University of Electronic Science and Technology of China*; Jiliang Wang, *Tsinghua University*;  
Chong Zhang, Hui Zheng, and Shen Yang, *University of Electronic Science and Technology of China*; Jinsong Han, *Zhejiang University*; Jian Li, *University of Electronic Science and Technology of China*

**Channel-Aware 5G RAN Slicing with Customizable Schedulers .....1767**  
Yongzhou Chen and Ruihao Yao, *UIUC*; Haitham Hassanieh, *EPFL*; Radhika Mittal, *UIUC*

**RF-CHORD: Towards Deployable RFID Localization System for Logistic Networks .....1783**  
Bo Liang, *Peking University and Alibaba Group*; Purui Wang, *Massachusetts Institute of Technology*; Renjie Zhao, *University of California San Diego*; Heyu Guo, *Peking University*; Pengyu Zhang and Junchen Guo, *Alibaba Group*;  
Shunmin Zhu, *Tsinghua University and Alibaba Group*; Hongqiang Harry Liu, *Alibaba Group*; Xinyu Zhang, *University of California San Diego*; Chenren Xu, *Peking University, Zhongguancun Laboratory, and Key Laboratory of High Confidence Software Technologies, Ministry of Education (PKU)*

**Exploring Practical Vulnerabilities of Machine Learning-based Wireless Systems ..... 1801**  
Zikun Liu, Changming Xu, and Emerson Sie, *University of Illinois Urbana-Champaign*; Gagandeep Singh, *University of Illinois Urbana-Champaign and VMware Research*; Deepak Vasish, *University of Illinois Urbana-Champaign*