

# **18th USENIX Symposium on Networked Systems Design and Implementation (NSDI'21)**

Online  
12 – 14 April 2021

Volume 1 of 2

ISBN: 978-1-7138-2906-5

**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© (2021) by Usenix Association  
All rights reserved.

Printed with permission by Curran Associates, Inc. (2021)

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)

**18th USENIX Symposium on  
Networked Systems Design and Implementation  
(NSDI '21)**

**April 12–14, 2021**

**Monday, April 12**

**Datacenter Networking and SDNs**

**Accessing Cloud with Disaggregated Software-Defined Router** . . . . . 1  
Hua Shao, *Tsinghua University*; Xiaoliang Wang, *Tencent and Nanjing University*; Yuanwei Lu, Yanbo Yu, and Shengli Zheng, *Tencent*; Youjian Zhao, *Tsinghua University*

**CodedBulk: Inter-Datacenter Bulk Transfers using Network Coding** . . . . . 15  
Shih-Hao Tseng, Saksham Agarwal, and Rachit Agarwal, *Cornell University*; Hitesh Ballani, *Microsoft Research*; Ao Tang, *Cornell University*

**Twenty Years After: Hierarchical Core-Stateless Fair Queueing** . . . . . 29  
Zhuolong Yu, Jingfeng Wu, and Vladimir Braverman, *Johns Hopkins University*; Ion Stoica, *UC Berkeley*; Xin Jin, *Peking University*

**Breaking the Transience-Equilibrium Nexus: A New Approach to Datacenter Packet Transport** . . . . . 47  
Shiyu Liu and Ahmad Ghalayini, *Stanford University*; Mohammad Alizadeh, *MIT*; Balaji Prabhakar and Mendel Rosenblum, *Stanford University*; Anirudh Sivaraman, *NYU*

**Running BGP in Data Centers at Scale** . . . . . 65  
Anubhavnidhi Abhashkumar and Kausik Subramanian, *University of Wisconsin–Madison*; Alexey Andreyev, Hyojeong Kim, Nanda Kishore Salem, Jingyi Yang, and Petr Lapukhov, *Facebook*; Aditya Akella, *University of Wisconsin–Madison*; Hongyi Zeng, *Facebook*

**Orion: Google’s Software-Defined Networking Control Plane** . . . . . 83  
Andrew D. Ferguson, Steve Gribble, Chi-Yao Hong, Charles Killian, Waqar Mohsin, Henrik Muehe, Joon Ong, Leon Poutievski, Arjun Singh, Lorenzo Vicisano, Richard Alimi, Shawn Shuoshuo Chen, Mike Conley, Subhasree Mandal, Karthik Nagaraj, Kondapa Naidu Bollineni, Amr Sabaa, Shidong Zhang, Min Zhu, and Amin Vahdat, *Google*

**Verification and Formal Methods**

**Metha: Network Verifiers Need To Be Correct Too!** . . . . . 99  
Rüdiger Birkner, Tobias Brodmann, Petar Tsankov, Laurent Vanbever, and Martin Vechev, *ETH Zürich*

**Finding Invariants of Distributed Systems: It’s a Small (Enough) World After All** . . . . . 115  
Travis Hance, Marijn Heule, Ruben Martins, and Bryan Parno, *Carnegie Mellon University*

**Avenir: Managing Data Plane Diversity with Control Plane Synthesis** . . . . . 133  
Eric Hayden Campbell, *Cornell*; William T. Hallahan, *Yale*; Priya Srikumar, *Cornell*; Carmelo Cascone, *ONF*; Jed Liu, *Intel*; Vignesh Ramamurthy, *Infosys*; Hossein Hojjat, *Tehran & TeIAS*; Ruzica Piskac and Robert Soulé, *Yale*; Nate Foster, *Cornell*

**Don’t Yank My Chain: Auditable NF Service Chaining** . . . . . 155  
Guyue Liu and Hugo Sadok, *Carnegie Mellon University*; Anne Kohlbrenner, *Princeton University*; Bryan Parno, Vyas Sekar, and Justine Sherry, *Carnegie Mellon University*

**Network Management**

**Contracting Wide-area Network Topologies to Solve Flow Problems Quickly** . . . . . 175  
Firas Abuzaid, *Microsoft Research and Stanford University*; Srikanth Kandula, Behnaz Arzani, and Ishai Menache, *Microsoft Research*; Matei Zaharia and Peter Bailis, *Stanford University*

**Cost-effective Cloud Edge Traffic Engineering with CASCARA** . . . . . 201  
Rachee Singh, Sharad Agarwal, Matt Calder, and Paramvir Bahl, *Microsoft*

**A Social Network Under Social Distancing: Risk-Driven Backbone Management During COVID-19 and Beyond . . . .217**  
Yiting Xia, *MPI-INF and Facebook*; Ying Zhang, *Facebook*; Zhizhen Zhong, *MIT and Facebook*; Guanqing Yan, Chiun Lin Lim, Satyajeeet Singh Ahuja, Soshant Bali, and Alexander Nikolaidis, *Facebook*; Kimia Ghobadi, *Johns Hopkins University*; Manya Ghobadi, *MIT*

**Staying Alive: Connection Path Reselection at the Edge . . . . . 233**  
Raul Landa, Lorenzo Saino, Lennert Buytenhek, and Joao Taveira Araujo, *Fastly*

**Debugging Transient Faults in Data Centers using Synchronized Network-wide Packet Histories . . . . . 253**  
Pravein Govindan Kannan, *IBM Research - India*; Nishant Budhdev, Raj Joshi, and Mun Choon Chan, *National University of Singapore*

## Web and Video

**Alohamora: Reviving HTTP/2 Push and Preload by Adapting Policies On the Fly . . . . . 269**  
Nikhil Kansal, Murali Ramanujam, and Ravi Netravali, *UCLA*

**Oblique: Accelerating Page Loads Using Symbolic Execution. . . . . 289**  
Ronny Ko and James Mickens, *Harvard University*; Blake Loring, *Royal Holloway, University of London*; Ravi Netravali, *UCLA*

**SENSEI: Aligning Video Streaming Quality with Dynamic User Sensitivity . . . . . 303**  
Xu Zhang and Yiyang Ou, *University of Chicago*; Siddhartha Sen, *Microsoft Research*; Junchen Jiang, *University of Chicago*

## Tuesday, April 13

### Databases and Analytics

**GAIA: A System for Interactive Analysis on Distributed Graphs Using a High-Level Language. . . . . 321**  
Zhengping Qian, Chenqiang Min, Longbin Lai, Yong Fang, Gaofeng Li, Youyang Yao, Bingqing Lyu, Xiaoli Zhou, Zhimin Chen, and Jingren Zhou, *Alibaba Group*

**TEGRA: Efficient Ad-Hoc Analytics on Evolving Graphs . . . . . 337**  
Anand Padmanabha Iyer, *Microsoft Research and University of California, Berkeley*; Qifan Pu, *Google*; Kishan Patel, *Two Sigma*; Joseph E. Gonzalez and Ion Stoica, *University of California, Berkeley*

**Unifying Timestamp with Transaction Ordering for MVCC with Decentralized Scalar Timestamp . . . . . 357**  
Xingda Wei, Rong Chen, Haibo Chen, Zhaoguo Wang, Zhenhan Gong, and Binyu Zang, *Shanghai Jiao Tong University*

**When to Hedge in Interactive Services. . . . . 373**  
Mia Primorac, *Oracle Labs*; Katerina Argyraki and Edouard Bugnion, *EPFL*

**Move Fast and Meet Deadlines: Fine-grained Real-time Stream Processing with Cameo. . . . . 389**  
Le Xu, *University of Illinois at Urbana-Champaign*; Shivaram Venkataraman, *UW-Madison*; Indranil Gupta, *University of Illinois at Urbana-Champaign*; Luo Mai, *University of Edinburgh*; Rahul Potharaju, *Microsoft*

**WHIZ: Data-Driven Analytics Execution. . . . . 407**  
Robert Grandl, *Google*; Arjun Singhvi, *University of Wisconsin–Madison*; Raajay Viswanathan, *Uber Technologies Inc.*; Aditya Akella, *University of Wisconsin–Madison*

### Mobile and IoT

**Pushing the Physical Limits of IoT Devices with Programmable Metasurfaces . . . . . 425**  
Lili Chen, *Northwest University (China) and University of Massachusetts Amherst*; Wenjun Hu, *Yale University*; Kyle Jamieson, *Princeton University*; Xiaojiang Chen and Dingyi Fang, *Northwest University (China)*; Jeremy Gummesson, *University of Massachusetts Amherst*

**Bootstrapping Battery-free Wireless Networks: Efficient Neighbor Discovery and Synchronization in the Face of Intermittency. . . . . 439**  
Kai Geissdoerfer and Marco Zimmerling, *TU Dresden*

**AIRCODE: Hidden Screen-Camera Communication on an Invisible and Inaudible Dual Channel . . . . . 457**  
Kun Qian, *Tsinghua University and University of California San Diego*; Yumeng Lu, Zheng Yang, Kai Zhang, Kehong Huang, and Xinjun Cai, *Tsinghua University*; Chenshu Wu, *University of Maryland College Park*; Yunhao Liu, *Tsinghua University and Michigan State University*

<b>Device-Based LTE Latency Reduction at the Application Layer</b> . . . . .	<b>471</b>
Zhaowei Tan and Jinghao Zhao, <i>UCLA</i> ; Yuanjie Li, <i>Tsinghua University</i> ; Yifei Xu, <i>Peking University</i> ; Songwu Lu, <i>UCLA</i>	
<b>System Performance and Programmability</b>	
<b>BMC: Accelerating Memcached using Safe In-kernel Caching and Pre-stack Processing</b> . . . . .	<b>487</b>
Yoann Ghigoff, <i>Orange Labs, Sorbonne Université, Inria, LIP6</i> ; Julien Sopena, <i>Sorbonne Université, LIP6</i> ; Kahina Lazri, <i>Orange Labs</i> ; Antoine Blin, <i>Gandi</i> ; Gilles Muller, <i>Inria</i>	
<b>Segcache: a memory-efficient and scalable in-memory key-value cache for small objects</b> . . . . .	<b>503</b>
Juncheng Yang, <i>Carnegie Mellon University</i> ; Yao Yue, <i>Twitter</i> ; Rashmi Vinayak, <i>Carnegie Mellon University</i>	
<b>When Cloud Storage Meets RDMA</b> . . . . .	<b>519</b>
Yixiao Gao, <i>Nanjing University and Alibaba Group</i> ; Qiang Li, Lingbo Tang, Yongqing Xi, Pengcheng Zhang, Wenwen Peng, Bo Li, Yaohui Wu, Shaozong Liu, Lei Yan, Fei Feng, Yan Zhuang, Fan Liu, Pan Liu, Xingkui Liu, Zhongjie Wu, Junping Wu, and Zheng Cao, <i>Alibaba Group</i> ; Chen Tian, <i>Nanjing University</i> ; Jinbo Wu, Jiayi Zhu, Haiyong Wang, Dennis Cai, and Jiasheng Wu, <i>Alibaba Group</i>	
<b>Prism: Proxies without the Pain</b> . . . . .	<b>535</b>
Yutaro Hayakawa, <i>LINE Corporation</i> ; Michio Honda, <i>University of Edinburgh</i> ; Douglas Santry, <i>Apple Inc.</i> ; Lars Eggert, <i>NetApp</i>	
<b>Programming Network Stack for Middleboxes with Rubik</b> . . . . .	<b>551</b>
Hao Li, <i>Xi'an Jiaotong University</i> ; Changhao Wu, <i>Xi'an Jiaotong University and Brown University</i> ; Guangda Sun, Peng Zhang, and Danfeng Shan, <i>Xi'an Jiaotong University</i> ; Tian Pan, <i>Beijing University of Posts and Telecommunications</i> ; Chengchen Hu, <i>Xilinx Labs Asia Pacific</i>	
<b>Flightplan: Dataplane Disaggregation and Placement for P4 Programs</b> . . . . .	<b>571</b>
Nik Sultana, John Sonchack, Hans Giesen, Isaac Pedisich, Zhaoyang Han, Nishanth Shyamkumar, Shivani Burad, André DeHon, and Boon Thau Loo, <i>University of Pennsylvania</i>	
<b>Distributed Systems</b>	
<b>MilliSort and MilliQuery: Large-Scale Data-Intensive Computing in Milliseconds</b> . . . . .	<b>593</b>
Yilong Li, <i>Stanford University</i> ; Seo Jin Park, <i>MIT CSAIL</i> ; John Ousterhout, <i>Stanford University</i>	
<b>EPaxos Revisited</b> . . . . .	<b>613</b>
Sarah Tollman, <i>Stanford University</i> ; Seo Jin Park, <i>MIT CSAIL</i> ; John Ousterhout, <i>Stanford University</i>	
<b>Ship Compute or Ship Data? Why Not Both?</b> . . . . .	<b>633</b>
Jie You, <i>University of Michigan</i> ; Jingfeng Wu, <i>Johns Hopkins University</i> ; Xin Jin, <i>Peking University</i> ; Mosharaf Chowdhury, <i>University of Michigan</i>	
<b>Caerus: NIMBLE Task Scheduling for Serverless Analytics</b> . . . . .	<b>653</b>
Hong Zhang, <i>UC Berkeley</i> ; Yupeng Tang and Anurag Khandelwal, <i>Yale University</i> ; Jingrong Chen, <i>Duke University</i> ; Ion Stoica, <i>UC Berkeley</i>	
<b>Ownership: A Distributed Futures System for Fine-Grained Tasks</b> . . . . .	<b>671</b>
Stephanie Wang, Eric Liang, and Edward Oakes, <i>UC Berkeley and Anyscale</i> ; Ben Hindman, Frank Sifei Luan, Audrey Cheng, and Ion Stoica, <i>UC Berkeley</i>	
<b>Fault-Tolerant Replication with Pull-Based Consensus in MongoDB</b> . . . . .	<b>687</b>
Siyuan Zhou, <i>MongoDB Inc.</i> ; Shuai Mu, <i>Stony Brook University</i>	
<b>Wednesday, April 14</b>	
<b>Machine Learning in a Systems Context</b>	
<b>Mistify: Automating DNN Model Porting for On-Device Inference at the Edge</b> . . . . .	<b>705</b>
Peizhen Guo, Bo Hu, and Wenjun Hu, <i>Yale University</i>	
<b>Elastic Resource Sharing for Distributed Deep Learning</b> . . . . .	<b>721</b>
Changho Hwang and Taehyun Kim, <i>KAIST</i> ; Sunghyun Kim, <i>MIT</i> ; Jinwoo Shin and Kyoungsoo Park, <i>KAIST</i>	

**ATP: In-network Aggregation for Multi-tenant Learning** . . . . . 741  
ChonLam Lao, *Tsinghua University*; Yanfang Le and Kshiteej Mahajan, *University of Wisconsin-Madison*; Yixi Chen and Wenfei Wu, *Tsinghua University*; Aditya Akella and Michael Swift, *University of Wisconsin-Madison*

**On the Use of ML for Blackbox System Performance Prediction** . . . . . 763  
Silvery Fu, *UC Berkeley*; Saurabh Gupta and Radhika Mittal, *UIUC*; Sylvia Ratnasamy, *UC Berkeley*

**Scaling Distributed Machine Learning with In-Network Aggregation**. . . . . 785  
Amedeo Sapia, Marco Canini, and Chen-Yu Ho, *KAUST*; Jacob Nelson, *Microsoft*; Panos Kalnis, *KAUST*; Changhoon Kim, *Barefoot Networks*; Arvind Krishnamurthy, *University of Washington*; Masoud Moshref, *Barefoot Networks*; Dan Ports, *Microsoft*; Peter Richtarik, *KAUST*

## Wireless Sensing

**Efficient Wideband Spectrum Sensing Using MEMS Acoustic Resonators**. . . . . 809  
Junfeng Guan, Jitian Zhang, Ruochen Lu, Hyungjoo Seo, Jin Zhou, Songbin Gong, and Haitham Hassanieh, *University of Illinois at Urbana-Champaign*

**WiForce: Wireless Sensing and Localization of Contact Forces on a Space Continuum** . . . . . 827  
Agrim Gupta, Cédric Girerd, Manideep Dunna, Qiming Zhang, Raghav Subbaraman, Tania Morimoto, and Dinesh Bharadia, *University of California, San Diego*

**MAVL: Multiresolution Analysis of Voice Localization** . . . . . 845  
Mei Wang, Wei Sun, and Lili Qiu, *The University of Texas at Austin*

**From Conception to Retirement: a Lifetime Story of a 3-Year-Old Wireless Beacon System in the Wild** . . . . . 859  
Yi Ding, *Alibaba Group, University of Minnesota*; Ling Liu, *Shanghai Jiao Tong University*; Yu Yang, *Rutgers University*; Yunhuai Liu, *Peking University*; Desheng Zhang, *Rutgers University*; Tian He, *Alibaba Group, University of Minnesota*

**EarFisher: Detecting Wireless Eavesdroppers by Stimulating and Sensing Memory EMR** . . . . . 873  
Cheng Shen, *Peking University*; Jun Huang, *Massachusetts Institute of Technology*

## Wireless

**Adapting Wireless Mesh Network Configuration from Simulation to Reality via Deep Learning based Domain Adaptation**. . . . . 887  
Junyang Shi and Mo Sha, *State University of New York at Binghamton*; Xi Peng, *University of Delaware*

**Practical Null Steering in Millimeter Wave Networks** . . . . . 903  
Sohrab Madani and Suraj Jog, *University of Illinois Urbana Champaign*; Jesus O. Lacruz and Joerg Widmer, *IMDEA Networks*; Haitham Hassanieh, *University of Illinois Urbana Champaign*

**SyncScatter: Enabling WiFi like synchronization and range for WiFi backscatter Communication** . . . . . 923  
Manideep Dunna, Miao Meng, Po-Han Wang, Chi Zhang, Patrick Mercier, and Dinesh Bharadia, *University of California, San Diego*

**Verification and Redesign of OFDM Backscatter** . . . . . 939  
Xin Liu, *University of Maryland, Baltimore County*; Zicheng Chi, *Cleveland State University*; Wei Wang, Yao Yao, Pei Hao, and Ting Zhu, *University of Maryland, Baltimore County*

**Simplifying Backscatter Deployment: Full-Duplex LoRa Backscatter** . . . . . 955  
Mohamad Katanbaf, *Jeeva Wireless and University of Washington*; Anthony Weinand and Vamsi Talla, *Jeeva Wireless*

**One Protocol to Rule Them All: Wireless Network-on-Chip using Deep Reinforcement Learning** . . . . . 973  
Suraj Jog, Zikun Liu, Antonio Franques, and Vimuth Fernando, *University of Illinois at Urbana Champaign*; Sergi Abadal, *Polytechnic University of Catalonia*; Josep Torrellas and Haitham Hassanieh, *University of Illinois at Urbana Champaign*

## Measurement

**LightGuardian: A Full-Visibility, Lightweight, In-band Telemetry System Using Sketchlets** . . . . . 991  
Yikai Zhao, Kaicheng Yang, and Zirui Liu, *Peking University*; Tong Yang, *Peking University and Peng Cheng Laboratory*; Li Chen, *Huawei Theory Lab*; Shiyi Liu, Naiqian Zheng, Ruixin Wang, and Hanbo Wu, *Peking University*; Yi Wang, *Southern University of Science and Technology and Peng Cheng Laboratory*; Nicholas Zhang, *Huawei Theory Lab*

**Fast and Light Bandwidth Testing for Internet Users** .....1011  
Xinlei Yang, Xianlong Wang, Zhenhua Li, and Yunhao Liu, *Tsinghua University*; Feng Qian, *University of Minnesota*;  
Liangyi Gong, *Tsinghua University*; Rui Miao, *Alibaba Group*; Tianyin Xu, *University of Illinois Urbana-Champaign*

**Toward Nearly-Zero-Error Sketching via Compressive Sensing**..... 1027  
Qun Huang, *Peking University and Pengcheng Lab*; Siyuan Sheng, *Institute of Computing Technology, CAS*; Xiang Chen,  
*Peking University and Pengcheng Lab and Fuzhou University*; Yungang Bao, *Institute of Computing Technology, CAS*;  
Rui Zhang, Yanwei Xu, and Gong Zhang, *Huawei Theory Department*