

22nd USENIX Symposium on Networked Systems Design and Implementation (NSDI'25)

Philadelphia, Pennsylvania, USA
28-30 April 2025

Volume 1 of 2

ISBN: 979-8-3313-1947-2

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© (2025) by Usenix Association
All rights reserved.

Printed with permission by Curran Associates, Inc. (2025)

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

22nd USENIX Symposium on Networked Systems Design and Implementation (NSDI '25)

April 28–30, 2025
Philadelphia, PA

Monday, April 28

Data Centers Queuing and Routing

PRED: Performance-oriented Random Early Detection for Consistently Stable Performance in Datacenters 1
Xinle Du, *Huawei Technologies*; Tong Li, *Renmin University of China*; Guangmeng Zhou, Zhuotao Liu, Hanlin Huang, and Xiangyu Gao, *Tsinghua University*; Mowei Wang and Kun Tan, *Huawei Technologies*; Ke Xu, *Tsinghua University*

Rajomon: Decentralized and Coordinated Overload Control for Latency-Sensitive Microservices 21
Jiali Xing, Akis Giannoukos, Paul Loh, Shuyue Wang, and Justin Qiu, *University of Pennsylvania*; Henri Maxime Demoulin, *DBOS, Inc*; Konstantinos Kallas, *University of California, Los Angeles*; Benjamin C. Lee, *University of Pennsylvania*

Learnings from Deploying Network QoS Alignment to Application Priorities for Storage Services 37
Matthew Buckley and Parsa Pazhooheshy, *Google and University of Toronto*; Z. Morley Mao, Nandita Dukkkipati, Hamid Hajabdolali Bazzaz, Priyaranjan Jha, Yingjie Bi, and Steve Middlekauff, *Google*; Yashar Ganjali, *University of Toronto*

D₁SC: Backpressure Mitigation In Multi-tier Applications With Distributed Shared Connection 55
Brice Ekane and Djob Mvondo, *Univ. Rennes, Inria, CNRS, IRISA, France*; Renaud Lachaize, *Univ. Grenoble Alpes, CNRS, Inria, Grenoble INP, LIG, 38000 Grenoble, France*; Yérom-David Bromberg, *Univ. Rennes, Inria, CNRS, IRISA, France*; Alain Tchana, *Univ. Grenoble Alpes, CNRS, Inria, Grenoble INP, LIG, 38000 Grenoble, France*; Daniel Hagimont, *IRIT, Université de Toulouse, CNRS, Toulouse INP, UT3 Toulouse, France*

Data Plane Programmability 1

Enabling Silent Telemetry Data Transmission with InvisiFlow 71
Yinda Zhang, *University of Pennsylvania*; Liangcheng Yu, *University of Pennsylvania and Microsoft Research*; Gianni Antichi, *Politecnico di Milano and Queen Mary University of London*; Ran Ben Basat, *University College London*; Vincent Liu, *University of Pennsylvania*

Unlocking ECMP Programmability for Precise Traffic Control 87
Yadong Liu, *Tencent*; Yunming Xiao, *University of Michigan*; Xuan Zhang, Weizhen Dang, Huihui Liu, Xiang Li, and Zekun He, *Tencent*; Jilong Wang, *Tsinghua University*; Aleksandar Kuzmanovic, *Northwestern University*; Ang Chen, *University of Michigan*; Congcong Miao, *Tencent*

Enabling Portable and High-Performance SmartNIC Programs with Alkali 107
Jiaxin Lin, *UT Austin*; Zhiyuan Guo, *UCSD*; Mihir Shah, *NVIDIA*; Tao Ji, *Microsoft*; Yiyang Zhang, *UCSD*; Daehyeok Kim and Aditya Akella, *UT Austin*

Scaling IP Lookup to Large Databases using the CRAM Lens 127
Robert Chang and Pradeep Dogga, *University of California, Los Angeles*; Andy Fingerhut, *Cisco Systems*; Victor Rios and George Varghese, *University of California, Los Angeles*

Data Center Resource Scheduling

Quicksand: Harnessing Stranded Datacenter Resources with Granular Computing 147
Zhenyuan Ruan, *MIT CSAIL*; Shihang Li, *Brown University*; Kaiyan Fan, *MIT CSAIL*; Seo Jin Park, *USC*; Marcos K. Aguilera, *VMware Research by Broadcom*; Adam Belay, *MIT CSAIL*; Malte Schwarzkopf, *Brown University*

Beehive: A Scalable Disaggregated Memory Runtime Exploiting Asynchrony of Multithreaded Programs 167
Quanxi Li, Hong Huang, Ying Liu, and Yanwen Xia, *Institute of Computing Technology, Chinese Academy of Sciences; University of Chinese Academy of Sciences*; Jie Zhang, *Peking University*; Mosong Zhou, *Huawei Cloud*; Xiaobing Feng and Huimin Cui, *Institute of Computing Technology, Chinese Academy of Sciences; University of Chinese Academy of Sciences*; Quan Chen, *Shanghai Jiao Tong University*; Yizhou Shan, *Huawei Cloud*; Chenxi Wang, *Institute of Computing Technology, Chinese Academy of Sciences; University of Chinese Academy of Sciences*

Making Serverless Pay-For-Use a Reality with Leopard	189
Tingjia Cao, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau, and Tyler Caraza-Harter, <i>University of Wisconsin-Madison</i>	
GRANNY: Granular Management of Compute-Intensive Applications in the Cloud	205
Carlos Segarra, Simon Shillaker, Guo Li, and Eleftheria Mappoura, <i>Imperial College London</i> ; Rodrigo Bruno, <i>INESC-ID, Instituto Superior Técnico, University of Lisbon</i> ; Lluís Vilanova and Peter Pietzuch, <i>Imperial College London</i>	
Verification 1	
On Temporal Verification of Stateful P4 Programs	219
Delong Zhang, Chong Ye, and Fei He, <i>School of Software, BNRist, Tsinghua University, Beijing 100084, China</i> ; <i>Key Laboratory for Information System Security, MoE, China</i>	
NDD: A Decision Diagram for Network Verification	237
Zechun Li, Peng Zhang, and Yichi Zhang, <i>Xi'an Jiaotong University</i> ; Hongkun Yang, <i>Google</i>	
Smart Casual Verification of the Confidential Consortium Framework	259
Heidi Howard, Markus A. Kuppe, Edward Ashton, and Amaury Chamayou, <i>Azure Research, Microsoft</i> ; Natacha Crooks, <i>Azure Research, Microsoft, and UC Berkeley</i>	
VEP: A Two-stage Verification Toolchain for Full eBPF Programmability	277
Xiwei Wu, Yueyang Feng, Tianyi Huang, Xiaoyang Lu, Shengkai Lin, Lihan Xie, Shizhen Zhao, and Qinxiong Cao, <i>Shanghai Jiao Tong University</i>	
Failure and Diagnosis	
MeshTest: End-to-End Testing for Service Mesh Traffic Management	301
Naiqian Zheng, Tianshuo Qiao, Xuanzhe Liu, and Xin Jin, <i>Peking University</i>	
Preventing Network Bottlenecks: Accelerating Datacenter Services with Hotspot-Aware Placement for Compute and Storage	317
Hamid Hajabdolali Bazzaz, Yingjie Bi, and Weiwu Pang, <i>Google</i> ; Minlan Yu, <i>Harvard University</i> ; Ramesh Govindan, <i>University of Southern California</i> ; Neal Cardwell, Nandita Dukkkipati, Meng-Jung Tsai, Chris DeForest, and Yuxue Jin, <i>Google</i> ; Charles Carver, <i>Columbia University</i> ; Jan Kopański, Liqun Cheng, and Amin Vahdat, <i>Google</i>	
Enhancing Network Failure Mitigation with Performance-Aware Ranking	335
Pooria Namyar and Arvin Ghavidel, <i>University of Southern California</i> ; Daniel Crankshaw, Daniel S. Berger, Kevin Hsieh, and Srikanth Kandula, <i>Microsoft</i> ; Ramesh Govindan, <i>University of Southern California</i> ; Behnaz Arzani, <i>Microsoft</i>	
One-Size-Fits-None: Understanding and Enhancing Slow-Fault Tolerance in Modern Distributed Systems	359
Ruiming Lu, <i>University of Michigan and Shanghai Jiao Tong University</i> ; Yunchi Lu and Yuxuan Jiang, <i>University of Michigan</i> ; Guangtao Xue, <i>Shanghai Jiao Tong University</i> ; Peng Huang, <i>University of Michigan</i>	
All Things Transport	
Pyrrha: Congestion-Root-Based Flow Control to Eliminate Head-of-Line Blocking in Datacenter	379
Kexin Liu, Zhaochen Zhang, Chang Liu, and Yizhi Wang, <i>Nanjing University</i> ; Vamsi Addanki and Stefan Schmid, <i>TU Berlin</i> ; Qingyue Wang, Wei Chen, Xiaoliang Wang, and Jiaqi Zheng, <i>Nanjing University</i> ; Wenhao Sun, Tao Wu, Ke Meng, Fei Chen, Weiguang Wang, and Bingyang Liu, <i>Huawei, China</i> ; Wanchun Dou, Guihai Chen, and Chen Tian, <i>Nanjing University</i>	
eTran: Extensible Kernel Transport with eBPF	407
Zhongjie Chen, <i>Tsinghua University</i> ; Qingkai Meng, <i>Nanjing University</i> ; ChonLam Lao, <i>Harvard University</i> ; Yifan Liu and Fengyuan Ren, <i>Tsinghua University</i> ; Minlan Yu, <i>Harvard University</i> ; Yang Zhou, <i>UC Berkeley and UC Davis</i>	
White-Boxing RDMA with Packet-Granular Software Control	427
Chenxingyu Zhao and Jaehong Min, <i>University of Washington</i> ; Ming Liu, <i>University of Wisconsin-Madison</i> ; Arvind Krishnamurthy, <i>University of Washington</i>	
SIRD: A Sender-Informed, Receiver-Driven Datacenter Transport Protocol	451
Konstantinos Prasopoulos, <i>EPFL</i> ; Ryan Kosta, <i>UCSD</i> ; Edouard Bugnion, <i>EPFL</i> ; Marios Kogias, <i>Imperial College London</i>	

LLM Training and Resilience

Accelerating Design Space Exploration for LLM Training Systems with Multi-experiment Parallel Simulation... 473
Fei Gui, *Tsinghua University, BNRist, and Tsinghua Shenzhen International Graduate School*; Kaihui Gao and Li Chen, *Zhongguancun Laboratory*; Dan Li, *Tsinghua University*; Vincent Liu, *University of Pennsylvania*; Ran Zhang and Hongbing Yang, *Zhongguancun Laboratory*; Dian Xiong, *Tsinghua University*

Optimizing RLHF Training for Large Language Models with Stage Fusion 489
Yinmin Zhong, Zili Zhang, Bingyang Wu, and Shengyu Liu, *School of Computer Science, Peking University*; Yukun Chen, Changyi Wan, Hanpeng Hu, Lei Xia, Ranchen Ming, and Yibo Zhu, *StepFun*; Xin Jin, *School of Computer Science, Peking University*

Minder: Faulty Machine Detection for Large-scale Distributed Model Training 505
Yangtao Deng, *Tsinghua University*; Xiang Shi and Zhuo Jiang, *ByteDance*; Xingjian Zhang, *Tsinghua University*; Lei Zhang, Zhang Zhang, Bo Li, Zuquan Song, Hang Zhu, and Gaohong Liu, *ByteDance*; Fuliang Li, *Northeastern University*; Shuguang Wang, Haibin Lin, and Jianxi Ye, *ByteDance*; Minlan Yu, *Harvard University*

Holmes: Localizing Irregularities in LLM Training with Mega-scale GPU Clusters 523
Zhiyi Yao and Pengbo Hu, *Fudan University and Tencent*; Congcong Miao and Xuya Jia, *Tencent*; Zuning Liang and Yuedong Xu, *Fudan University*; Chunzhi He, Hao Lu, Mingzhuo Chen, Xiang Li, Zekun He, Yachen Wang, and Xianneng Zou, *Tencent*; Junchen Jiang, *University of Chicago*

SimAI: Unifying Architecture Design and Performance Tuning for Large-Scale Large Language Model Training with Scalability and Precision 541
Xizheng Wang, *Alibaba Cloud and Tsinghua University*; Qingxu Li, Yichi Xu, and Gang Lu, *Alibaba Cloud*; Dan Li, *Tsinghua University*; Li Chen, *Zhongguancun Laboratory*; Heyang Zhou, *Alibaba Cloud*; Linkang Zheng, *Alibaba Cloud and South China University of Technology*; Sen Zhang, Yikai Zhu, Yang Liu, Pengcheng Zhang, Kun Qian, Kunling He, Jiaqi Gao, Ennan Zhai, Dennis Cai, and Binzhang Fu, *Alibaba Cloud*

ByteCheckpoint: A Unified Checkpointing System for Large Foundation Model Development 559
Borui Wan, *The University of Hong Kong*; Mingji Han, Yiyao Sheng, Yanghua Peng, Haibin Lin, Mofan Zhang, Zhichao Lai, Menghan Yu, Junda Zhang, Zuquan Song, and Xin Liu, *ByteDance Inc.*; Chuan Wu, *The University of Hong Kong*

Video and Cloud Gaming

Mowgli: Passively Learned Rate Control for Real-Time Video 579
Neil Agarwal and Rui Pan, *Princeton University*; Francis Y. Yan, *University of Illinois Urbana-Champaign*; Ravi Netravali, *Princeton University*

Dissecting and Streamlining the Interactive Loop of Mobile Cloud Gaming 595
Yang Li, Jiaying Qiu, Hongyi Wang, and Zhenhua Li, *Tsinghua University*; Feng Qian, *University of Southern California*; Jing Yang, *Tsinghua University*; Hao Lin, *Tsinghua University and UIUC*; Yunhao Liu, *Tsinghua University*; Bo Xiao and Xiaokang Qin, *Ant Group*; Tianyin Xu, *UIUC*

Region-based Content Enhancement for Efficient Video Analytics at the Edge 613
Weijun Wang, *Institute for AI Industry Research (AIR), Tsinghua University*; Liang Mi, Shaowei Cen, and Haipeng Dai, *State Key Laboratory for Novel Software Technology, Nanjing University*; Yuanchun Li, *Institute for AI Industry Research (AIR), Tsinghua University*; Xiaoming Fu, *University of Göttingen*; Yunxin Liu, *Institute for AI Industry Research (AIR), Tsinghua University*

Tooth: Toward Optimal Balance of Video QoE and Redundancy Cost by Fine-Grained FEC in Cloud Gaming Streaming 635
Congkai An, Huanhuan Zhang, Shibo Wang, Jingyang Kang, Anfu Zhou, Liang Liu and Huadong Ma, *Beijing University of Posts and Telecommunications*; Zili Meng, *Hong Kong University of Science and Technology*; Delei Ma, Yusheng Dong, and Xiaogang Lei, *Well-Link Times Inc.*

AsTree: An Audio Subscription Architecture Enabling Massive-Scale Multi-Party Conferencing 653
Tong Meng, Wenfeng Li, Chao Yuan, Changqing Yan, and Le Zhang, *ByteDance Inc.*

Tuesday, April 29

Infra For ML

AutoCCL: Automated Collective Communication Tuning for Accelerating Distributed and Parallel DNN Training 667

Guanbin Xu, Zhihao Le, Yinhe Chen, Zhiqi Lin, and Zewen Jin, *University of Science and Technology of China*; Youshan Miao, *Microsoft Research*; Cheng Li, *University of Science and Technology of China*; Anhui Province Key Laboratory of Biomedical Imaging and Intelligent Processing; Institute of Artificial Intelligence, Hefei Comprehensive National Science Center

OPTiREDUCE: Resilient and Tail-Optimal AllReduce for Distributed Deep Learning in the Cloud 685

Ertza Warraich, *Purdue University*; Omer Shabtai and Khalid Manaa, *Nvidia*; Shay Vargaftik, *VMware Research*; Yonatan Piasetzky and Matty Kadosh, *Nvidia*; Lalith Suresh, *Feldera*; Muhammad Shahbaz, *University of Michigan*

Efficient Direct-Connect Topologies for Collective Communications 705

Liangyu Zhao, *University of Washington*; Siddharth Pal, *RTX BBN Technologies*; Tapan Chugh, *University of Washington*; Weiyang Wang, *MIT CSAIL*; Jason Fantl, Prithwish Basu, and Joud Khoury, *RTX BBN Technologies*; Arvind Krishnamurthy, *University of Washington*

SuperServe: Fine-Grained Inference Serving for Unpredictable Workloads. 739

Alind Khare and Dhruv Garg, *Georgia Tech*; Sukrit Kalra, *UC Berkeley*; Snigdha Grandhi, *Adobe*; Ion Stoica, *UC Berkeley*; Alexey Tumanov, *Georgia Tech*

Fast Scalable Consensus

Pineapple: Unifying Multi-Paxos and Atomic Shared Registers 759

Tigran Bantikyan, *Northwestern*; Jonathan Zarnstorff, *unaffiliated*; Te-Yen Chou, *CMU*; Lewis Tseng, *UMass Lowell*; Roberto Palmieri, *Lehigh University*

Ladder: A Convergence-based Structured DAG Blockchain for High Throughput and Low Latency 779

Dengcheng Hu, Jianrong Wang, Xiulong Liu, and Hao Xu, *Tianjin University*; Xujing Wu, *Jd.Com, Inc*; Muhammad Shahzad, *North Carolina State University*; Guyue Liu, *Peking University*; Keqiu Li, *Tianjin University*

Vegeta: Enabling Parallel Smart Contract Execution in Leaderless Blockchains 795

Tianjing Xu and Yongqi Zhong, *Shanghai Jiao Tong University*; Yiming Zhang, *Shanghai Jiao Tong University and Shanghai Key Laboratory of Trusted Data Circulation, Governance and Web3*; Ruofan Xiong, *Xiamen University*; Jingjing Zhang, *Fudan University*; Guangtao Xue and Shengyun Liu, *Shanghai Jiao Tong University and Shanghai Key Laboratory of Trusted Data Circulation, Governance and Web3*

Shoal++: High Throughput DAG BFT Can Be Fast and Robust! 813

Balaji Arun and Zekun Li, *Aptos Labs*; Florian Suri-Payer, *Cornell University*; Sourav Das, *UIUC*; Alexander Spiegelman, *Aptos Labs*

Operational Experiences

Learning Production-Optimized Congestion Control Selection for Alibaba Cloud CDN 827

Xuan Zeng, *Alibaba Cloud*; Haoran Xu, *Sun Yat-sen University*; Chen Chen and Xumiao Zhang, *Alibaba Cloud*; Xiaoxi Zhang and Xu Chen, *Sun Yat-sen University*; Guihai Chen, *Nanjing University*; Yubing Qiu, Yiping Zhang, Chong Hao, and Ennan Zhai, *Alibaba Cloud*

GPU-Disaggregated Serving for Deep Learning Recommendation Models at Scale 847

Lingyun Yang, *Hong Kong University of Science and Technology*; Yongchen Wang and Yinghao Yu, *Alibaba Group*; Qizhen Weng, *Hong Kong University of Science and Technology*; Jianbo Dong, Kan Liu, Chi Zhang, Yanyi Zi, Hao Li, Zechao Zhang, Nan Wang, Yu Dong, Menglei Zheng, Lanlan Xi, Xiaowei Lu, Liang Ye, Guodong Yang, Binzhang Fu, Tao Lan, Liping Zhang, and Lin Qu, *Alibaba Group*; Wei Wang, *Hong Kong University of Science and Technology*

Evolution of Aegis: Fault Diagnosis for AI Model Training Service in Production. 865

Jianbo Dong, Kun Qian, Pengcheng Zhang, Zhilong Zheng, Liang Chen, Fei Feng, Yichi Xu, Yikai Zhu, Gang Lu, Xue Li, Zhihui Ren, Zhicheng Wang, Bin Luo, Peng Zhang, Yang Liu, Yanqing Chen, Yu Guan, Weicheng Wang, Chaojie Yang, Yang Zhang, Man Yuan, Hanyu Zhao, Yong Li, Zihan Zhao, Shan Li, Xianlong Zeng, Zhiping Yao, Binzhang Fu, Ennan Zhai, Wei Lin, Chao Wang, and Dennis Cai, *Alibaba Cloud*

PAPAYA Federated Analytics Stack: Engineering Privacy, Scalability and Practicality	883
Harish Srinivas, Graham Cormode, Mehrdad Honarkhah, Samuel Lurye, Jonathan Hehir, Lunwen He, George Hong, Ahmed Magdy, Dzmitry Huba, Kaikai Wang, Shen Guo, and Shoubhik Bhattacharya, <i>Meta</i>	
Middleboxes	
HA/TCP: A Reliable and Scalable Framework for TCP Network Functions	899
Haoyu Gu, Ali José Mashtizadeh, and Bernard Wong, <i>University of Waterloo</i>	
High-level Programming for Application Networks	915
Xiangfeng Zhu, Yuyao Wang, Banruo Liu, Yongtong Wu, and Nikola Bojanic, <i>University of Washington</i> ; Jingrong Chen, <i>Duke University</i> ; Gilbert Louis Bernstein and Arvind Krishnamurthy, <i>University of Washington</i> ; Sam Kumar, <i>University of Washington and UCLA</i> ; Ratul Mahajan, <i>University of Washington</i> ; Danyang Zhuo, <i>Duke University</i>	
State-Compute Replication: Parallelizing High-Speed Stateful Packet Processing	937
Qiongwen Xu, <i>Rutgers University</i> ; Sebastiano Miano, <i>Politecnico di Milano</i> ; Xiangyu Gao and Tao Wang, <i>New York University</i> ; Adithya Murugadass and Songyuan Zhang, <i>Rutgers University</i> ; Anirudh Sivaraman, <i>New York University</i> ; Gianni Antichi, <i>Queen Mary University of London and Politecnico di Milano</i> ; Srinivas Narayana, <i>Rutgers University</i>	
MTP: Transport for In-Network Computing	959
Tao Ji, <i>UT Austin</i> ; Rohan Vardekar and Balajee Vamanan, <i>University of Illinois Chicago</i> ; Brent E. Stephens, <i>Google and University of Utah</i> ; Aditya Akella, <i>UT Austin</i>	
Rethinking Data Center Efficiency	
ONCache: A Cache-Based Low-Overhead Container Overlay Network	979
Shengkai Lin, Shizhen Zhao, Peirui Cao, and Xinchu Han, <i>Shanghai Jiao Tong University</i> ; Quan Tian, Wenfeng Liu, Qi Wu, and Donghai Han, <i>Broadcom</i> ; Xinbing Wang, <i>Shanghai Jiao Tong University</i>	
GREEN: Carbon-efficient Resource Scheduling for Machine Learning Clusters	999
Kaiqiang Xu and Decang Sun, <i>iSING Lab, Hong Kong University of Science and Technology</i> ; Han Tian, <i>USTC</i> ; Junxue Zhang and Kai Chen, <i>iSING Lab, Hong Kong University of Science and Technology</i>	
The Benefits and Limitations of User Interrupts for Preemptive Userspace Scheduling	1015
Linsong Guo, Danial Zuberi, Tal Garfinkel, and Amy Ousterhout, <i>UC San Diego</i>	
Securing Public Cloud Networks with Efficient Role-based Micro-Segmentation	1033
Sathya Kumaran Mani and Kevin Hsieh, <i>Microsoft</i> ; Santiago Segarra, <i>Rice University</i> ; Ranveer Chandra, <i>Microsoft</i> ; Yajie Zhou, <i>University of Maryland</i> ; Srikanth Kandula, <i>Microsoft</i>	
RDMA	
Mitigating Scalability Walls of RDMA-based Container Networks	1049
Wei Liu, <i>Tsinghua University and Alibaba Cloud</i> ; Kun Qian, <i>Alibaba Cloud</i> ; Zhenhua Li, <i>Tsinghua University</i> ; Feng Qian, <i>University of Southern California</i> ; Tianyin Xu, <i>UIUC</i> ; Yunhao Liu, <i>Tsinghua University</i> ; Yu Guan, Shuhong Zhu, Hongfei Xu, Lanlan Xi, Chao Qin, and Ennan Zhai, <i>Alibaba Cloud</i>	
Eden: Developer-Friendly Application-Integrated Far Memory	1067
Anil Yelam, Stewart Grant, and Saarth Deshpande, <i>UC San Diego</i> ; Nadav Amit, <i>Technion, Israel Institute of Technology</i> ; Radhika Niranjana Mysore, <i>VMware Research Group</i> ; Amy Ousterhout, <i>UC San Diego</i> ; Marcos K. Aguilera, <i>VMware Research Group</i> ; Alex C. Snoeren, <i>UC San Diego</i>	
Achieving Wire-Latency Storage Systems by Exploiting Hardware ACKs	1085
Qing Wang, Ji Wu Shu, Jing Wang, and Yuhao Zhang, <i>Tsinghua University</i>	
ODRP: On-Demand Remote Paging with Programmable RDMA	1101
Zixuan Wang, Xingda Wei, Jinyu Gu, Hongrui Xie, Rong Chen, and Haibo Chen, <i>Institute of Parallel and Distributed Systems, SEIEE, Shanghai Jiao Tong University</i>	

Storage

- Understanding and Profiling NVMe-over-TCP Using ntprof** 1117
Yuyuan Kang and Ming Liu, *University of Wisconsin-Madison*
- Building an Elastic Block Storage over EBOFs Using Shadow Views** 1137
Sheng Jiang, *Carnegie Mellon University*; Ming Liu, *University of Wisconsin-Madison*
- Pushing the Limits of In-Network Caching for Key-Value Stores** 1155
Gyuyeong Kim, *Sungshin Women's University*

Cellular and Wireless

- CellReplay: Towards accurate record-and-replay for cellular networks** 1169
William Sentosa, *University of Illinois Urbana-Champaign*; Balakrishnan Chandrasekaran, *VU Amsterdam*;
P. Brighten Godfrey, *University of Illinois Urbana-Champaign and Broadcom*; Haitham Hassanieh, *EPFL*
- Large Network UWB Localization: Algorithms and Implementation** 1187
Nakul Garg and Irtaza Shahid, *University of Maryland, College Park*; Ramanujan K Sheshadri, *Nokia Bell Labs*;
Karthikeyan Sundaresan, *Georgia Institute of Technology*; Nirupam Roy, *University of Maryland, College Park*
- Towards Energy Efficient 5G vRAN Servers** 1205
Anuj Kalia, *Microsoft*; Nikita Lazarev, *MIT*; Leyang Xue, *University of Edinburgh*; Xenofon Foukas
and Bozidar Radunovic, *Microsoft*; Francis Y. Yan, *Microsoft Research and UIUC*
- Building Massive MIMO Baseband Processing on a Single-Node Supercomputer** 1221
Xincheng Xie, Wentao Hou, Zerui Guo, and Ming Liu, *University of Wisconsin-Madison*
- Efficient Multi-WAN Transport for 5G with OTTER** 1243
Mary Hogan, *Oberlin College*; Gerry Wan, *Google*; Yiming Qiu, *University of Michigan*; Sharad Agarwal
and Ryan Beckett, *Microsoft*; Rachee Singh, *Cornell University*; Paramvir Bahl, *Microsoft*

Wednesday, April 30

Verification 2

- Verifying maximum link loads in a changing world** 1269
Tibor Schneider, *ETH Zürich*; Stefano Vissicchio, *University College London*; Laurent Vanbever, *ETH Zürich*
- A Layered Formal Methods Approach to Answering Queue-related Queries** 1289
Divya Raghunathan, Maria Apostolaki, and Aarti Gupta, *Princeton University*
- Runtime Protocol Refinement Checking for Distributed Protocol Implementations** 1305
Ding Ding, Zhanghan Wang, Jinyang Li, and Aurojit Panda, *NYU*
- CEGS: Configuration Example Generalizing Synthesizer** 1327
Jianmin Liu, *Tsinghua University*; Li Chen, *Zhongguancun Laboratory*; Dan Li, *Tsinghua University*;
Yukai Miao, *Zhongguancun Laboratory*

Security

- Suppressing BGP Zombies with Route Status Transparency** 1349
Yosef Edery Anahory, *The Hebrew University of Jerusalem*; Jie Kong, Nicholas Scaglione, and Justin Furuness,
University of Connecticut; Hemi Leibowitz, *The College of Management Academic Studies*; Amir Herzberg and
Bing Wang, *University of Connecticut*; Yossi Gilad, *The Hebrew University of Jerusalem*
- ValidaTor: Domain Validation over Tor** 1367
Jens Frieß, *National Research Center for Applied Cybersecurity ATHENE and Technische Universität Darmstadt*;
Haya Schulmann, *National Research Center for Applied Cybersecurity ATHENE and Goethe-Universität Frankfurt*;
Michael Waidner, *National Research Center for Applied Cybersecurity ATHENE and Technische Universität Darmstadt*

From Address Blocks to Authorized Prefixes: Redesigning RPKI ROV with a Hierarchical Hashing Scheme for Fast and Memory-Efficient Validation 1381
 Zedong Ni, *Computer Network Information Center, Chinese Academy of Sciences; and School of Cyber Science & Engineering, Southeast University*; Yinbo Xu, Hui Zou, and Yanbiao Li, *Computer Network Information Center, Chinese Academy of Sciences; and University of Chinese Academy of Sciences*; Guang Cheng, *School of Cyber Science & Engineering, Southeast University*; and Purple Mountain Laboratories; Gaogang Xie, *Computer Network Information Center, Chinese Academy of Sciences; and University of Chinese Academy of Sciences*

PreAcher: Secure and Practical Password Pre-Authentication by Content Delivery Networks 1399
 Shihan Lin, *Duke University*; Suting Chen, *Northwestern University*; Yunming Xiao, *University of Michigan*; Yanqi Gu, *University of California, Irvine*; Aleksandar Kuzmanovic, *Northwestern University*; Xiaowei Yang, *Duke University*

Data Plane Programmability 2

ClubHeap: A High-Speed and Scalable Priority Queue for Programmable Packet Scheduling 1421
 Zhikang Chen, *Tsinghua University*; Haoyu Song, *Futurewei Technologies*; Zhiyu Zhang and Yang Xu, *Fudan University*; Bin Liu, *Tsinghua University*

Self-Clocked Round-Robin Packet Scheduling 1437
 Erfan Sharafzadeh, *Johns Hopkins University and Hewlett Packard Labs*; Raymond Matson, *University of California Riverside*; Jean Tourrilhes and Puneet Sharma, *Hewlett Packard Labs*; Soudeh Ghorbani, *Johns Hopkins University and Meta*

Everything Matters in Programmable Packet Scheduling 1467
 Albert Gran Alcoz, *ETH Zürich*; Balázs Vass, *BME-TMIT*; Pooria Namyar, *USC*; Behnaz Arzani, *Microsoft Research*; Gábor Rétvári, *BME-TMIT*; Laurent Vanbever, *ETH Zürich*

When P4 Meets Run-to-completion Architecture 1487
 Hao Zheng, *State Key Laboratory for Novel Software Technology, Nanjing University, China*; Xin Yan, *Huawei, China*; Wenbo Li, Jiaqi Zheng, and Xiaoliang Wang, *State Key Laboratory for Novel Software Technology, Nanjing University, China*; Qingqing Zhao, Luyou He, Xiaofei Lai, Feng Gao, and Fuguang Huang, *Huawei, China*; Wanchun Dou, Guihai Chen, and Chen Tian, *State Key Laboratory for Novel Software Technology, Nanjing University, China*

ML for Networks

Mutant: Learning Congestion Control from Existing Protocols via Online Reinforcement Learning 1507
 Lorenzo Pappone, *Computer Science Department, Saint Louis University*; Alessio Sacco, *DAUIN, Politecnico di Torino*; Flavio Esposito, *Computer Science Department, Saint Louis University*

CATO: End-to-End Optimization of ML-Based Traffic Analysis Pipelines 1523
 Gerry Wan, *Stanford University*; Shinan Liu, *University of Chicago*; Francesco Bronzino, *ENS Lyon*; Nick Feamster, *University of Chicago*; Zakir Durumeric, *Stanford University*

Resolving Packets from Counters: Enabling Multi-scale Network Traffic Super Resolution via Composable Large Traffic Model 1541
 Xizheng Wang, *Tsinghua University and Zhongguancun Laboratory*; Libin Liu and Li Chen, *Zhongguancun Laboratory*; Dan Li, *Tsinghua University*; Yukai Miao and Yu Bai, *Zhongguancun Laboratory*

BFTBrain: Adaptive BFT Consensus with Reinforcement Learning 1563
 Chenyuan Wu and Haoyun Qin, *University of Pennsylvania*; Mohammad Javad Amiri, *Stony Brook University*; Boon Thau Loo, *University of Pennsylvania*; Dahlia Malkhi, *UC Santa Barbara*; Ryan Marcus, *University of Pennsylvania*