

17th USENIX Symposium on Operating Systems Design and Implementation (OSDI'23)

Boston, Massachusetts, USA
10-12 July 2023

Volume 1 of 2

ISBN: 978-1-7138-8396-8

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. (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

17th USENIX Symposium on Operating Systems Design and Implementation (OSDI '23)

July 10–12, 2023
Boston, MA, USA

Monday, July 10

Make Your Bits Go Faster

Ship your Critical Section, Not Your Data: Enabling Transparent Delegation with TCLocks 1
Vishal Gupta, *EPFL*; Kumar Kartikeya Dwivedi, *SRMIST*; Yugesh Kothari, Yueyang Pan, Diyu Zhou, and Sanidhya Kashyap, *EPFL*

RON: One-Way Circular Shortest Routing to Achieve Efficient and Bounded-waiting Spinlocks17
Shiwu Lo, Han-Ting Lin, Yao-Hung Hsieh, and Chao-Ting Lin, *National Chung Cheng University*; Yu-Hsueh Fang, *National Cheng Kung University*; Ching-Shen Lin, *National Chung Cheng University*; Ching-Chun (Jim) Huang, *National Cheng Kung University*; Kam Yiu Lam, *City University of Hong Kong*; Yuan-Hao Chang, *Academia Sinica, Taiwan*

Userspace Bypass: Accelerating Syscall-intensive Applications..... 33
Zhe Zhou, Yanxiang Bi, Junpeng Wan, and Yangfan Zhou, *Fudan University*; Zhou Li, *University of California, Irvine*

Triangulating Python Performance Issues with SCALENE..... 51
Emery D. Berger, Sam Stern, and Juan Altmayer Pizzorno, *University of Massachusetts Amherst*

Relational Debugging — Pinpointing Root Causes of Performance Problems 65
Xiang (Jenny) Ren, Sitao Wang, Zhuqi Jin, David Lion, and Adrian Chiu, *University of Toronto*; Tianyin Xu, *University of Illinois at Urbana-Champaign*; Ding Yuan, *University of Toronto*

Secure Your Bits I

Accountable authentication with privacy protection: The Larch system for universal login 81
Emma Dauterman, *UC Berkeley*; Danny Lin, *Woodinville High School*; Henry Corrigan-Gibbs, *MIT*; David Mazières, *Stanford University*

K9db: Privacy-Compliant Storage For Web Applications By Construction..... 99
Kinan Dak Albab, Ishan Sharma, Justus Adam, Benjamin Kilimnik, Aaron Jeyaraj, Raj Paul, Artem Agvastian, Leonhard Spiegelberg, and Malte Schwarzkopf, *Brown University*

Encrypted Databases Made Secure Yet Maintainable117
Mingyu Li, *Shanghai Jiao Tong University*; Shanghai AI Laboratory; Engineering Research Center for Domain-specific Operating Systems, Ministry of Education, China; Xuyang Zhao and Le Chen, *Shanghai Jiao Tong University*; Engineering Research Center for Domain-specific Operating Systems, Ministry of Education, China; Cheng Tan, *Northeastern University*; Huorong Li and Sheng Wang, *Alibaba Group*; Zeyu Mi, *Shanghai Jiao Tong University*; Engineering Research Center for Domain-specific Operating Systems, Ministry of Education, China; Yubin Xia, *Shanghai Jiao Tong University*; Shanghai AI Laboratory; Engineering Research Center for Domain-specific Operating Systems, Ministry of Education, China; Feifei Li, *Alibaba Group*; Haibo Chen, *Shanghai Jiao Tong University*; Engineering Research Center for Domain-specific Operating Systems, Ministry of Education, China

LVMT: An Efficient Authenticated Storage for Blockchain 135
Chenxing Li, *Shanghai Tree-Graph Blockchain Research Institute*; Sidi Mohamed Beillahi, *University of Toronto*; Guang Yang and Ming Wu, *Shanghai Tree-Graph Blockchain Research Institute*; Wei Xu, *Tsinghua University*; Fan Long, *University of Toronto*

Honeycomb: Secure and Efficient GPU Executions via Static Validation..... 155
Haohui Mai, *PrivacyCore Inc.*; Jiacheng Zhao, *SKLP, Institute of Computing Technology, CAS*; Zhongguancun Laboratory; and UCAS; Hongren Zheng, *IIS, Tsinghua University*; Yiyang Zhao, *SKLP, Institute of Computing Technology, CAS*; and UCAS; Zibin Liu, *BUPT*; Mingyu Gao, *IIS, Tsinghua University*; Cong Wang, *IDEA Shenzhen*; Huimin Cui, *SKLP, Institute of Computing Technology, CAS*; and UCAS; Xiaobing Feng, *SKLP, Institute of Computing Technology, CAS*; Zhongguancun Laboratory; and UCAS; Christos Kozyrakis, *PrivacyCore Inc. and Stanford*

Secure Your Bits II

An Extensible Orchestration and Protection Framework for Confidential Cloud Computing173
Adil Ahmad and Alex Schultz, *Arizona State University*; Byoungyoung Lee, *Seoul National University*; Pedro Fonseca, *Purdue University*

Nimble: Rollback Protection for Confidential Cloud Services..... 193
Sebastian Angel, *Microsoft Research*; Aditya Basu, *Penn State University*; Weidong Cui, *Microsoft Research*;
Trent Jaeger, *Penn State University*; Stella Lau, *MIT CSAIL*; Srinath Setty, *Microsoft Research*; Sudheesh Singanamalla, *University of Washington*

Kerveros: Efficient and Scalable Cloud Admission Control..... 209
Sultan Mahmud Sajal, *Microsoft Research and Pennsylvania State University*; Luke Marshall and Beibin Li, *Microsoft Research*; Shandan Zhou and Abhisek Pan, *Microsoft Azure*; Konstantina Mellou and Deepak Narayanan, *Microsoft Research*; Timothy Zhu, *Pennsylvania State University*; David Dion and Thomas Moscibroda, *Microsoft Azure*;
Ishai Menache, *Microsoft Research*

Security and Performance in the Delegated User-level Virtualization 227
Jiahao Chen, *Institute of Parallel and Distributed Systems, SEIEE, Shanghai Jiao Tong University*; Engineering Research Center for Domain-specific Operating Systems, Ministry of Education, China; Dingji Li, *Institute of Parallel and Distributed Systems, SEIEE, Shanghai Jiao Tong University*; Engineering Research Center for Domain-specific Operating Systems, Ministry of Education, China; MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University; Zeyu Mi, Yuxuan Liu, and Binyu Zang, *Institute of Parallel and Distributed Systems, SEIEE, Shanghai Jiao Tong University*; Engineering Research Center for Domain-specific Operating Systems, Ministry of Education, China; Haibing Guan, *Shanghai Key Laboratory of Scalable Computing and Systems, Shanghai Jiao Tong University*; Haibo Chen, *Institute of Parallel and Distributed Systems, SEIEE, Shanghai Jiao Tong University*; Engineering Research Center for Domain-specific Operating Systems, Ministry of Education, China

Core slicing: closing the gap between leaky confidential VMs and bare-metal cloud..... 247
Ziqiao Zhou, *Microsoft Research*; Yizhou Shan, *University of California, San Diego*; Weidong Cui, Xinyang Ge, Marcus Peinado, and Andrew Baumann, *Microsoft Research*

Tuesday, July 11

Expanding, Hardening, and Deploying Your Bits

ExoFlow: A Universal Workflow System for Exactly-Once DAGs..... 269
Siyuan Zhuang, *UC Berkeley*; Stephanie Wang, *UC Berkeley and Anyscale*; Eric Liang and Yi Cheng, *Anyscale*;
Ion Stoica, *UC Berkeley*

Hyrax: Fail-in-Place Server Operation in Cloud Platforms..... 287
Jialun Lyu, *Microsoft Azure and University of Toronto*; Marisa You, Celine Irvine, Mark Jung, Tyler Narmore, Jacob Shapiro, Luke Marshall, and Savyasachi Samal, *Microsoft Azure*; Ioannis Manousakis and Lisa Hsu, *Formerly of Microsoft Azure*; Preetha Subbarayalu, Ashish Raniwala, Brijesh Warriar, and Ricardo Bianchini, *Microsoft Azure*;
Bianca Schroeder, *University of Toronto*; Daniel S. Berger, *Microsoft Azure and University of Washington*

NCC: Natural Concurrency Control for Strictly Serializable Datastores by Avoiding the Timestamp-Inversion Pitfall 305
Haonan Lu, *University at Buffalo*; Shuai Mu, *Stony Brook University*; Siddhartha Sen, *Microsoft Research*;
Wyatt Lloyd, *Princeton University*

Conveyor: One-Tool-Fits-All Continuous Software Deployment at Meta..... 325
Boris Grubic, *Meta*; Yang Wang, *Meta and the Ohio State University*; Tyler Petrochko, Ran Yaniv, Brad Jones, David Callies, Matt Clarke-Lauer, and Dan Kelley, *Meta*; Soteris Demetriou, *Meta and Imperial College London*; Kenny Yu and Chunqiang Tang, *Meta*

Query Your Bits

Chardonnay: Fast and General Datacenter Transactions for On-Disk Databases 343
Tamer Eldeeb and Xincheng Xie, *Columbia University*; Philip A. Bernstein, *Microsoft Research*; Asaf Cidon and Junfeng Yang, *Columbia University*

ScaleDB: A Scalable, Asynchronous In-Memory Database 361
Syed Akbar Mehdi, *The University of Texas at Austin*; Deukyeon Hwang and Simon Peter, *University of Washington*;
Lorenzo Alvisi, *Cornell University*

VBASE: Unifying Online Vector Similarity Search and Relational Queries via Relaxed Monotonicity 377
Qianxi Zhang, Shuotao Xu, Qi Chen, and Guoxin Sui, *Microsoft Research Asia*; Jiadong Xie, *Microsoft Research Asia and East China Normal University*; Zhizhen Cai and Yaoqi Chen, *Microsoft Research Asia and University of Science and Technology of China*; Yinxuan He, *Microsoft Research Asia and Renmin University of China*; Yuqing Yang, Fan Yang, Mao Yang, and Lidong Zhou, *Microsoft Research Asia*

Detecting Transactional Bugs in Database Engines via Graph-Based Oracle Construction 397
Zu-Ming Jiang and Si Liu, *ETH Zurich*; Manuel Rigger, *National University of Singapore*; Zhendong Su, *ETH Zurich*

Take Out the TraChe: Maximizing (Tra)nsactional Ca(che) Hit Rate 419
Audrey Cheng, David Chu, Terrance Li, Jason Chan, Natacha Crooks, Joseph M. Hellerstein, and Ion Stoica, *UC Berkeley*;
Xiangyao Yu, *University of Wisconsin—Madison*

Store Your Bits

Replicating Persistent Memory Key-Value Stores with Efficient RDMA Abstraction 441
Qing Wang, Youyou Lu, Jing Wang, and Jiwu Shu, *Tsinghua University*

eZNS: An Elastic Zoned Namespace for Commodity ZNS SSDs 461
Jiahong Min and Chenxingyu Zhao, *University of Washington*; Ming Liu, *University of Wisconsin-Madison*;
Arvind Krishnamurthy, *University of Washington*

SEPH: Scalable, Efficient, and Predictable Hashing on Persistent Memory 479
Chao Wang, Junliang Hu, Tsun-Yu Yang, Yuhong Liang, and Ming-Chang Yang, *The Chinese University of Hong Kong*

No Provisioned Concurrency: Fast RDMA-codedigned Remote Fork for Serverless Computing 497
Xingda Wei, *Institute of Parallel and Distributed Systems, SEITEE, Shanghai Jiao Tong University, and Shanghai AI Laboratory*; Fangming Lu, Tianxia Wang, Jinyu Gu, and Yuhang Yang, *Institute of Parallel and Distributed Systems, SEITEE, Shanghai Jiao Tong University*; Rong Chen, *Institute of Parallel and Distributed Systems, SEITEE, Shanghai Jiao Tong University, and Shanghai AI Laboratory*; Haibo Chen, *Institute of Parallel and Distributed Systems, SEITEE, Shanghai Jiao Tong University*

Manage Your Bits I

Johnny Cache: the End of DRAM Cache Conflicts (in Tiered Main Memory Systems)..... 519
Baptiste Lepers, *Université de Neuchâtel*; Willy Zwaenepoel, *University of Sydney*

TAILCHECK: A Lightweight Heap Overflow Detection Mechanism with Page Protection and Tagged Pointers..... 535
Amogha Udupa Shankaranarayana Gopal, Raveendra Soori, Michael Ferdman, and Dongyoon Lee, *Stony Brook University*

SMART: A High-Performance Adaptive Radix Tree for Disaggregated Memory 553
Xuchuan Luo, *School of Computer Science, Fudan University*; Pengfei Zuo, *Huawei Cloud*; Jiacheng Shen and Jiazhen Gu, *The Chinese University of Hong Kong*; Xin Wang, *School of Computer Science, Fudan University*; Shanghai Key Laboratory of Intelligent Information Processing, Shanghai, China; Michael R. Lyu, *The Chinese University of Hong Kong*; Yangfan Zhou, *School of Computer Science, Fudan University*; Shanghai Key Laboratory of Intelligent Information Processing, Shanghai, China

ORC: Increasing Cloud Memory Density via Object Reuse with Capabilities 573
Vasily A. Sartakov, Lluís Vilanova, and Munir Geden, *Imperial College London*; David Eysers, *University of Otago*;
Takahiro Shinagawa, *The University of Tokyo*; Peter Pietzuch, *Imperial College London*

Manage Your Bits II

Global Capacity Management With Flux..... 589
Marius Eriksen, Kaushik Veeraraghavan, Yusuf Abdulghani, Andrew Birchall, Po-Yen Chou, Richard Cornew, Adela Kabiljo, Ranjith Kumar S, Maroo Lieuw, Justin Meza, Scott Michelson, Thomas Rohloff, Hayley Russell, Jeff Qin, and Chunqiang Tang, *Meta*

Defcon: Preventing Overload with Graceful Feature Degradation	607
Justin J. Meza, Thote Gowda, Ahmed Eid, Tomiwa Ijaware, Dmitry Chernyshev, Yi Yu, Md Nazim Uddin, Rohan Das, Chad Nachiappan, Sari Tran, Shuyang Shi, Tina Luo, David Ke Hong, Sankaralingam Panneerselvam, Hans Ragas, Svetlin Manavski, Weidong Wang, and Francois Richard, <i>Meta Platforms, Inc.</i>	
Cilantro: Performance-Aware Resource Allocation for General Objectives via Online Feedback	623
Romil Bhardwaj, <i>UC Berkeley</i> ; Kirthevasan Kandasamy, <i>University of Wisconsin-Madison</i> ; Asim Biswal, Wenshuo Guo, Benjamin Hindman, Joseph Gonzalez, Michael Jordan, and Ion Stoica, <i>UC Berkeley</i>	
Karma: Resource Allocation for Dynamic Demands	645
Midhul Vuppapapati, Giannis Fikioris, and Rachit Agarwal, <i>Cornell University</i> ; Asaf Cidon, <i>Columbia University</i> ; Anurag Khandelwal, <i>Yale University</i> ; Éva Tardos, <i>Cornell University</i>	

Wednesday, July 12

Train Your Bits I

AlpaServe: Statistical Multiplexing with Model Parallelism for Deep Learning Serving	663
Zhuohan Li and Lianmin Zheng, <i>UC Berkeley</i> ; Yinmin Zhong, <i>Peking University</i> ; Vincent Liu, <i>University of Pennsylvania</i> ; Ying Sheng, <i>Stanford University</i> ; Xin Jin, <i>Peking University</i> ; Yanping Huang and Zhifeng Chen, <i>Google</i> ; Hao Zhang, <i>UC San Diego</i> ; Joseph E. Gonzalez and Ion Stoica, <i>UC Berkeley</i>	
Cocktailer: Analyzing and Optimizing Dynamic Control Flow in Deep Learning	681
Chen Zhang, <i>Tsinghua University</i> ; Lingxiao Ma and Jilong Xue, <i>Microsoft Research</i> ; Yining Shi, <i>Peking University & Microsoft Research</i> ; Ziming Miao and Fan Yang, <i>Microsoft Research</i> ; Jidong Zhai, <i>Tsinghua University</i> ; Zhi Yang, <i>Peking University</i> ; Mao Yang, <i>Microsoft Research</i>	
Welder: Scheduling Deep Learning Memory Access via Tile-graph.	701
Yining Shi, <i>Peking University & Microsoft Research</i> ; Zhi Yang, <i>Peking University</i> ; Jilong Xue, Lingxiao Ma, Yuqing Xia, Ziming Miao, Yuxiao Guo, Fan Yang, and Lidong Zhou, <i>Microsoft Research</i>	
Effectively Scheduling Computational Graphs of Deep Neural Networks toward Their Domain-Specific Accelerators	719
Jie Zhao, <i>Information Engineering University</i> ; Siyuan Feng, <i>Shanghai Jiao Tong University</i> ; Xiaoqiang Dan, Fei Liu, Chengke Wang, Sheng Yuan, Wenyan Lv, and Qikai Xie, <i>Stream Computing Inc.</i>	

Train Your Bits II

EINet: Optimizing Tensor Programs with Derivation-Based Transformations	739
Liyan Zheng, Haojie Wang, Jidong Zhai, Muyan Hu, Zixuan Ma, Tuowei Wang, and Shuhong Huang, <i>Tsinghua University</i> ; Xupeng Miao, <i>Carnegie Mellon University</i> ; Shizhi Tang and Kezhao Huang, <i>Tsinghua University</i> ; Zhihao Jia, <i>Carnegie Mellon University</i>	
Hydro: Surrogate-Based Hyperparameter Tuning Service in Datacenters	757
Qinghao Hu, <i>Nanyang Technological University, S-Lab, NTU, and Shanghai AI Laboratory</i> ; Zhisheng Ye, <i>Shanghai AI Laboratory and Peking University</i> ; Meng Zhang, <i>Nanyang Technological University, S-Lab, NTU, and Shanghai AI Laboratory</i> ; Qiaoling Chen, <i>Shanghai AI Laboratory and National University of Singapore</i> ; Peng Sun, <i>Shanghai AI Laboratory and SenseTime Research</i> ; Yonggang Wen and Tianwei Zhang, <i>Nanyang Technological University</i>	
MGG: Accelerating Graph Neural Networks with Fine-Grained Intra-Kernel Communication-Computation Pipelining on Multi-GPU Platforms	779
Yuke Wang, Boyuan Feng, and Zheng Wang, <i>University of California Santa Barbara</i> ; Tong Geng, <i>University of Rochester</i> ; Kevin Barker and Ang Li, <i>Pacific Northwest National Laboratory</i> ; Yufei Ding, <i>University of California Santa Barbara</i>	
Optimizing Dynamic Neural Networks with Brainstorm	797
Weihao Cui, <i>Shanghai Jiao Tong University</i> ; Zhenhua Han, <i>Microsoft Research Asia</i> ; Lingji Ouyang, <i>University of Science and Technology of China</i> ; Yichuan Wang, <i>Shanghai Jiao Tong University</i> ; Ningxin Zheng, Lingxiao Ma, Yuqing Yang, Fan Yang, Jilong Xue, Lili Qiu, and Lidong Zhou, <i>Microsoft Research Asia</i> ; Quan Chen, <i>Shanghai Jiao Tong University</i> ; Haisheng Tan, <i>University of Science and Technology of China</i> ; Minyi Guo, <i>Shanghai Jiao Tong University</i>	

AdaEmbed: Adaptive Embedding for Large-Scale Recommendation Models 817
 Fan Lai, *University of Michigan*; Wei Zhang, Rui Liu, William Tsai, Xiaohan Wei, Yuxi Hu, Sabin Devkota, Jianyu Huang, Jongsoo Park, Xing Liu, Zeliang Chen, Ellie Wen, Paul Rivera, Jie You, and Chun-cheng Jason Chen, *Meta*; Mosharaf Chowdhury, *University of Michigan*

Verify Your Bits

BWoS: Formally Verified Block-based Work Stealing for Parallel Processing 833
 Jiawei Wang, *Huawei Dresden Research Center, Huawei Central Software Institute, Technische Universität Dresden*; Bohdan Trach, Ming Fu, Diogo Behrens, Jonathan Schwender, Yutao Liu, and Jitang Lei, *Huawei Dresden Research Center, Huawei Central Software Institute*; Viktor Vafeiadis, *MPI-SWS*; Hermann Härtig, *Technische Universität Dresden*; Haibo Chen, *Huawei Central Software Institute, Shanghai Jiao Tong University*

Spoq: Scaling Machine-Checkable Systems Verification in Coq 851
 Xupeng Li, Xuheng Li, Wei Qiang, Ronghui Gu, and Jason Nieh, *Columbia University*

Verifying vMVCC, a high-performance transaction library using multi-version concurrency control 871
 Yun-Sheng Chang, *MIT CSAIL*; Ralf Jung, *ETH Zurich*; Upamanyu Sharma, *MIT CSAIL*; Joseph Tassarotti, *New York University*; M. Frans Kaashoek and Nickolai Zeldovich, *MIT CSAIL*

Automated Verification of Idempotence for Stateful Serverless Applications 887
 Haoran Ding, Zhaoguo Wang, and Zhuohao Shen, *Institute of Parallel and Distributed Systems, SEITE, Shanghai Jiao Tong University*; Rong Chen, *Institute of Parallel and Distributed Systems, SEITE, Shanghai Jiao Tong University*; Shanghai AI Laboratory; Haibo Chen, *Institute of Parallel and Distributed Systems, SEITE, Shanghai Jiao Tong University*

Sharding the State Machine: Automated Modular Reasoning for Complex Concurrent Systems 911
 Travis Hance and Yi Zhou, *Carnegie Mellon University*; Andrea Lattuada, *VMware Research*; Reto Achermann, *University of British Columbia*; Alex Conway, *VMware Research*; Ryan Stutsman, *VMware Research and University of Utah*; Gerd Zellweger, *VMware Research*; Chris Hawblitzel, *Microsoft Research*; Jon Howell, *VMware Research*; Bryan Parno, *Carnegie Mellon University*

Transfer Your Bits

Flor: An Open High Performance RDMA Framework Over Heterogeneous RNICs 931
 Qiang Li, *Alibaba Group*; Yixiao Gao and Xiaoliang Wang, *Nanjing University*; Haonan Qiu, *Alibaba Group*; Yanfang Le, *AMD*; Derui Liu, *Alibaba Group*; Qiao Xiang, *Xiamen University*; Fei Feng, Peng Zhang, Bo Li, Jianbo Dong, Lingbo Tang, Hongqiang Harry Liu, Shaozong Liu, Weijie Li, Rui Miao, Yaohui Wu, Zhiwu Wu, Chao Han, Lei Yan, Zheng Cao, and Zhongjie Wu, *Alibaba Group*; Chen Tian and Guihai Chen, *Nanjing University*; Dennis Cai, Jinbo Wu, Jiaji Zhu and Jiesheng Wu, *Alibaba Group*; Jiwu Shu, *Xiamen University*

ShRing: Networking with Shared Receive Rings 949
 Boris Pismenny, *Technion & NVIDIA*; Adam Morrison, *Tel Aviv University*; Dan Tsafir, *Technion & VMware Research*

ServiceRouter: Hyperscale and Minimal Cost Service Mesh at Meta 969
 Harshit Saokar, *Meta*; Soteris Demetriou, *Meta and Imperial College London*; Nick Magerko, Max Kontorovich, Josh Kirstein, and Margot Leibold, *Meta*; Dimitrios Skarlatos, *Meta and Carnegie Mellon University*; Hitesh Khandelwal and Chunqiang Tang, *Meta*

Characterizing Off-path SmartNIC for Accelerating Distributed Systems 987
 Xingda Wei and Rongxin Cheng, *Institute of Parallel and Distributed Systems, SEITE, Shanghai Jiao Tong University, and Shanghai AI Laboratory*; Yuhan Yang, *Institute of Parallel and Distributed Systems, SEITE, Shanghai Jiao Tong University*; Rong Chen, *Institute of Parallel and Distributed Systems, SEITE, Shanghai Jiao Tong University, and Shanghai AI Laboratory*; Haibo Chen, *Institute of Parallel and Distributed Systems, SEITE, Shanghai Jiao Tong University*

Ensō : A Streaming Interface for NIC-Application Communication 1005
 Hugo Sadok and Nirav Atre, *Carnegie Mellon University*; Zhipeng Zhao, *Microsoft*; Daniel S. Berger, *Microsoft Research and University of Washington*; James C. Hoe, *Carnegie Mellon University*; Aurojit Panda, *New York University*; Justine Sherry, *Carnegie Mellon University*; Ren Wang, *Intel*