

USENIX ATC '25: 2025 USENIX Annual Technical Conference

July 7–9, 2025
Boston, MA, USA

Monday, July 7

Cloud Computing: Speed, Scale, and Serverless

Fast ACS: Low-Latency File-Based Ordered Message Delivery at Scale 1
Sushant Kumar Gupta, Anil Raghunath Iyer, Chang Yu, Neel Bagora, Olivier Pomerleau, Vivek Kumar, and Prunthaban Kanthakumar, *Google LLC*

Poby: SmartNIC-accelerated Image Provisioning for Coldstart in Clouds 19
Zihao Chang and Jiaqi Zhu, *SKLP, Institute of Computing Technology, CAS; University of Chinese Academy of Sciences*; Haifeng Sun, *Peking University*; Yunlong Xie, Kan Shi, Ninghui Sun, Yungang Bao, and Sa Wang, *SKLP, Institute of Computing Technology, CAS; University of Chinese Academy of Sciences*

Burst Computing: Quick, Sudden, Massively Parallel Processing on Serverless Resources 39
Daniel Barcelona-Pons, *Universitat Rovira i Virgili and Barcelona Supercomputing Center*; Aitor Arjona, Pedro García-López, Enrique Molina-Giménez, and Stepan Klymonchuk, *Universitat Rovira i Virgili*

DEEPSERVE: Serverless Large Language Model Serving at Scale 57
Junhao Hu, *Peking University and Key Lab of HCST (PKU), MOE*; Jiang Xu, Zhixia Liu, Yulong He, Yuetao Chen, Hao Xu, Jiang Liu, Jie Meng, Baoquan Zhang, Shining Wan, Gengyuan Dan, Zhiyu Dong, Zhihao Ren, and Changhong Liu, *Huawei Cloud*; Tao Xie, *Key Lab of HCST (PKU), MOE and Peking University*; Dayun Lin, Qin Zhang, Yue Yu, Hao Feng, Xusheng Chen, and Yizhou Shan, *Huawei Cloud*

Cosmic: Cost-Effective Support for Cloud-Assisted 3D Printing 73
Yuan Yao, *University of Southern California*; Chuan He and Chinedum Okwudire, *University of Michigan*; Harsha V. Madhyastha, *University of Southern California*

Accelerating ML Training: Parallelism, Tuning, and Modalities

GMI-DRL: Empowering Multi-GPU DRL with Adaptive-Grained Parallelism 89
Yuke Wang, *Rice University*; Boyuan Feng, *University of California Santa Barbara*; Zheng Wang, *University of California San Diego*; Guyue Huang, *University of California Santa Barbara*; Tong (Tony) Geng, *University of Rochester*; Ang Li, *Pacific Northwest National Laboratory*; Yufei Ding, *University of California San Diego*

mTuner: Accelerating Parameter-Efficient Fine-Tuning on Multi-GPU Servers with Elastic Tensor 105
Kezhao Huang, Siqi Zhu, Mingshu Zhai, Liyan Zheng, Kinman Lei, Jiaao He, Yuyang Jin, and Jidong Zhai, *Tsinghua University*

JENGA: Enhancing LLM Long-Context Fine-tuning with Contextual Token Sparsity 123
Tuowei Wang and Xingyu Chen, *Tsinghua University*; Kun Li and Ting Cao, *Microsoft Research*; Ju Ren and Yaoxue Zhang, *Tsinghua University*

FlexPipe: Maximizing Training Efficiency for Transformer-based Models with Variable-Length Inputs 143
Hairui Zhao, *Jilin University and University of California, Riverside*; Qi Tian, and Hongliang Li, *Jilin University and Key Laboratory of Symbolic Computation and Knowledge Engineering of the Ministry of Education, China*; Zizhong Chen, *University of California, Riverside*

Optimus: Accelerating Large-Scale Multi-Modal LLM Training by Bubble Exploitation 161
Weiqi Feng, *Harvard University*; Yangrui Chen, *ByteDance*; Shaoyu Wang, *University of Southern California*; Yanghua Peng and Haibin Lin, *ByteDance*; Minlan Yu, *Harvard University*

Networking: From Cloud to In-Network Intelligence

Towards Optimal Rack-scale μ s-level CPU Scheduling through In-Network Workload Shaping 179
Xudong Liao, *Hong Kong University of Science and Technology*; Han Tian, *University of Science and Technology of China*; Xinchun Wan, *Hong Kong University of Science and Technology*; Chaoliang Zeng, *BitIntelligence*; Hao Wang, *Hong Kong University of Science and Technology*; Junxue Zhang, *University of Science and Technology of China*; Mengyu Ma, *Inspur*; Guyue (Grace) Liu, *Peking University*; Kai Chen, *Hong Kong University of Science and Technology*

TGW: Operating an Efficient and Resilient Cloud Gateway at Scale 199
Yifan Yang, Lin He, Jiasheng Zhou, Xiaoyi Shi, Yichi Xu, and Shicheng Wang, *Tsinghua University*; Jinlong E, *Renmin University of China*; Ying Liu, *Tsinghua University*; Junwei Zhang, Zhuang Yuan, and Hengyang Xu, *Tencent*

MARC: Motion-Aware Rate Control for Mobile E-commerce Cloud Rendering 217
Yuankang Zhao, *Alibaba Group*; *University of Chinese Academy of Sciences*; Furong Yang, *Alibaba Group*; Gerui Lv, *University of Chinese Academy of Sciences*; Qinghua Wu, *University of Chinese Academy of Sciences and Purple Mountain Laboratories*; Yanmei Liu, Jiuhai Zhang, Yutang Peng, Feng Peng, Hongyu Guo, and Ying Chen, *Alibaba Group*; Zhenyu Li, *University of Chinese Academy of Sciences and Purple Mountain Laboratories*; Gaogang Xie, *University of Chinese Academy of Sciences and Computer Network Information Center, Chinese Academy of Sciences*

Accelerating Distributed Graph Learning by Using Collaborative In-Network Multicast and Aggregation 233
Zhaoyi Li, *Central South University and Nanyang Technological University*; Jiawei Huang, Yijun Li, and Jingling Liu, *Central South University*; Junxue Zhang, *Hong Kong University of Science and Technology*; Hui Li, Xiaojun Zhu, Shengwen Zhou, Jing Shao, Xiaojuan Lu, Qichen Su, and Jianxin Wang, *Central South University*; Chee Wei Tan, *Nanyang Technological University*; Yong Cui, *Tsinghua University*; Kai Chen, *Hong Kong University of Science and Technology*

Opening Up Kernel-Bypass TCP Stacks 249
Shinichi Awamoto and Michio Honda, *University of Edinburgh*

Operating Systems: Scheduling, Security, and Extensibility

GPREEMPT: GPU Preemptive Scheduling Made General and Efficient 263
Ruwen Fan and Tingxu Ren, *Tsinghua University*; Minhui Xie, *Renmin University of China*; Shiwei Gao, Jiwu Shu, and Youyou Lu, *Tsinghua University*

μ EFI: A Microkernel-Style UEFI with Isolation and Transparency 273
Le Chen, Yiyang Wu, Jinyu Gu, Yubin Xia, and Haibo Chen, *Shanghai Jiao Tong University*

PageFlex: Flexible and Efficient User-space Delegation of Linux Paging Policies with eBPF 291
Anil Yelam and Kan Wu, *Google*; Zhiyuan Guo, *UC San Diego*; Suli Yang, *Google*; Rajath Shashidhara, *University of Washington*; Wei Xu and Stanko Novaković, *Google*; Alex C. Snoeren, *Google and UC San Diego*; Kimberly Keeton, *Google*

ASTERINAS: A Linux ABI-Compatible, Rust-Based Framekernel OS with a Small and Sound TCB 307
Yuke Peng, *SUSTech*; Hongliang Tian, *Ant Group*; Junyang Zhang and Ruihan Li, *Peking University and Zhongguancun Laboratory*; Chengjun Chen and Jianfeng Jiang, *Ant Group*; Jinyi Xian, *SUSTech*; Xiaolin Wang, Chenren Xu, Diyu Zhou, and Yingwei Luo, *Peking University and Zhongguancun Laboratory*; Shoumeng Yan, *Ant Group*; Yinqian Zhang, *SUSTech*

Rex: Closing the language-verifier gap with safe and usable kernel extensions 325
Jinghao Jia and Ruowen Qin, *University of Illinois Urbana-Champaign*; Milo Craun and Egor Lukiyarov, *Virginia Tech*; Ayush Bansal and Minh Phan, *University of Illinois Urbana-Champaign*; Michael V. Le, Hubertus Franke, and Hani Jamjoom, *IBM*; Tianyin Xu, *University of Illinois Urbana-Champaign*; Dan Williams, *Virginia Tech*

The Programmable Data Plane: SmartNICs and Beyond

Barre: Empowering Simplified and Versatile Programmable Congestion Control in High-Speed AI Clusters 343
Yajuan Peng, *Shanghai Key Laboratory for Intelligence Information Processing, Fudan University, China*; Haoran Wei, Xiaolong Zhong, Junkai Huang, Haohan Xu, Zicheng Wang, Yang Bai, Zhuo Jiang, and Jianxi Ye, *ByteDance*; Xiaoliang Wang; Xiaoming Fu, *Shanghai Key Laboratory for Intelligence Information Processing, Fudan University, China*; Huichen Dai, *ByteDance*

FLB: Fine-grained Load Balancing for Lossless Datacenter Networks 365
Jinbin Hu, *Central South University, Hong Kong University of Science and Technology, Changsha University of Science and Technology*; Wenxue Li, Xiangzhou Liu, Junfeng Wang, and Bowen Liu, *Hong Kong University of Science and Technology*; Ping Yin, *Inspur*; Jianxin Wang and Jiawei Huang, *Central South University*; Kai Chen, *Hong Kong University of Science and Technology*

SNARY: A High-Performance and Generic SmartNIC-accelerated Retrieval System 381
Qiaoyin Gan, *Institute of Computing Technology, Chinese Academy of Sciences*; Heng Pan, *Computer Network Information Center, Chinese Academy of Sciences*; Luyang Li, Kai Lv, and Hongtao Guan, *Institute of Computing Technology, Chinese Academy of Sciences*; Zhaohua Wang, *Computer Network Information Center, Chinese Academy of Sciences*; Zhenyu Li, *Institute of Computing Technology, Chinese Academy of Sciences*; Gaogang Xie, *Computer Network Information Center, Chinese Academy of Sciences*

Minos : A Lightweight and Dynamic Defense against Traffic Analysis in Programmable Data Planes. 399
Zihao Wang, *Pengcheng Laboratory and Tsinghua Shenzhen International Graduate School*; Qing Li, Guorui Xie, Dan Zhao, Kejun Li, and Zhuochen Fan, *Pengcheng Laboratory*; Lianbo Ma, *Northeastern University*; Yong Jiang, *Pengcheng Laboratory and Tsinghua Shenzhen International Graduate School*

Performance: Benchmarking, Caching, and Workload Characterization

GeneralSparse: Bridging the Gap in SpMM for Pruned Large Language Model Inference on GPUs. 417
Yaoyu Wang, Xiao Guo, Junmin Xiao, De Chen, and Guangming Tan, *SKLP, Institute of Computing Technology, CAS; and University of Chinese Academy of Sciences*

HyCache: Hybrid Caching for Accelerating DNN Input Preprocessing Pipelines. 433
Keshav Vinayak Jha, *Independent Researcher*; Shweta Pandey, *Indian Institute of Science*; Murali Annavaram, *University of Southern California*; Arkaprava Basu, *Indian Institute of Science*

The Koala Benchmarks for the Shell: Characterization and Implications. 449
Evangelos Lamprou and Ethan Williams, *Brown University*; Georgios Kaoukis, *National Technical University of Athens*; Zhuoxuan Zhang, *Brown University*; Michael Greenberg, *Stevens Institute of Technology*; Konstantinos Kallas, *University of California, Los Angeles*; Lukas Lazarek and Nikos Vasilakis, *Brown University*

KVCache Cache in the Wild: Characterizing and Optimizing KVCache Cache at a Large Cloud Provider. 465
Jiahao Wang, Jinbo Han, and Xingda Wei, *Shanghai Jiao Tong University*; Sijie Shen, *Alibaba Group*; Dingyan Zhang, *Shanghai Jiao Tong University*; Chenguang Fang, *Alibaba Group*; Rong Chen, *Shanghai Jiao Tong University*; Wenjuan Yu, *Alibaba Group*; Haibo Chen, *Shanghai Jiao Tong University*

Tuesday, July 8

Storage Innovations: Logs, Tiers, and Modern Flash

LogCrisp: Fast Aggregated Analysis on Large-scale Compressed Logs by Enabling Two-Phase Pattern Extraction and Vectorized Queries. 483
Junyu Wei, Guangyan Zhang, and Junchao Chen, *Tsinghua University*; Qi Zhou, *Alibaba Cloud*

HotRAP: Hot Record Retention and Promotion for LSM-trees with Tiered Storage. 497
Jiansheng Qiu, Fangzhou Yuan, *Tsinghua University*; Mingyu Gao, and Huanchen Zhang, *Tsinghua University and Shanghai Qi Zhi Institute*

Mitigating Resource Usage Dependency in Sorting-based KV Stores on Hybrid Storage Devices via Operation Decoupling. 513
Qingyang Zhang and Yongkun Li, *University of Science and Technology of China*; Yubiao Pan, *Huaqiao University*; Haoting Tang, *University of Science and Technology of China*; Yinlong Xu, *University of Science and Technology of China, and Anhui Provincial Key Laboratory of High Performance Computing*

SolFS: An Operation-Log Versioning File System for Hash-free Efficient Mobile Cloud Backup 531
Riwei Pan, *Department of Computer Science, City University of Hong Kong*; Yu Liang, *ETH Zurich*; Lei Li and Hongchao Du, *Department of Computer Science, City University of Hong Kong*; Tei-Wei Kuo, *Delta Electronics and National Taiwan University*; Chun Jason Xue, *Mohamed bin Zayed University of Artificial Intelligence*

Z-LFS: A Zoned Namespace-tailored Log-structured File System for Commodity Small-zone ZNS SSDs. 547
Inhwi Hwang, *Seoul National University*; Sangjin Lee, *Chung-Ang University*; Sunggon Kim, *Seoul National University of Science and Technology*; Hyeonsang Eom, *Seoul National University*; Yongseok Son, *Chung-Ang University*

Serving Intelligence: Efficient LLM Inference at the Edge and Cloud

CLONE: Customizing LLMs for Efficient Latency-Aware Inference at the Edge 563
Chunlin Tian, Xinpeng Qin, Kahou Tam, Li Li, Zijian Wang, Yuanzhe Zhao, Minglei Zhang, and Chengzhong Xu, *University of Macau*

WEAVER: Efficient Multi-LLM Serving with Attention Offloading 587
Shiwei Gao, Qing Wang, Shaoxun Zeng, Youyou Lu, and Jiwu Shu, *Tsinghua University*

Torpor: GPU-Enabled Serverless Computing for Low-Latency, Resource-Efficient Inference 597
Minchen Yu, *The Chinese University of Hong Kong, Shenzhen*; and *Hong Kong University of Science and Technology*;
Ao Wang, *Alibaba Group*; Dong Chen, Haoxuan Yu, Xiaonan Luo, Zhuohao Li, and Wei Wang, *Hong Kong University of Science and Technology*; Ruichuan Chen, *Nokia Bell Labs*; Dapeng Nie, Haoran Yang, and Yu Ding, *Alibaba Group*

TOPPINGS: CPU-Assisted, Rank-Aware Adapter Serving for LLM Inference 613
Suyi Li, Hanfeng Lu, and Tianyuan Wu, *Hong Kong University of Science and Technology*; Minchen Yu, *The Chinese University of Hong Kong, Shenzhen*; Qizhen Weng, *TeleAI, China Telecom*; Xusheng Chen and Yizhou Shan, *Huawei Cloud*; Binhang Yuan and Wei Wang, *Hong Kong University of Science and Technology*

QFactory: Accelerating Quantized Large Language Model Serving with Qtile Graphs 631
Qihao Zhang, Mingshu Zhai, Rui Sun, and Jidong Zhai, *Tsinghua University*

Optimizing ML Execution: Compilers, Pipelines, and Runtimes

PluS: Highly Efficient and Expandable ML Compiler with Pluggable Graph Schedules 647
Ruofan Wu, *Renmin University of China*; Zhen Zheng, *Microsoft*; Feng Zhang, *Renmin University of China*;
Chuanjie Liu, *Microsoft*; Zaifeng Pan, *Renmin University of China*; Jidong Zhai, *Tsinghua University*; Xiaoyong Du, *Renmin University of China*

Obscura: Concealing Recomputation Overhead in Training of Large Language Models with Bubble-filling Pipeline Transformation 665
Yuzhou Huang and Yapeng Jiang, *Sun Yat-sen University*; Zicong Hong, *Hong Kong University of Science and Technology*; Wuhui Chen, *Sun Yat-sen University*; Bin Wang and Weixi Zhu, *Huawei Technologies*; Yue Yu, *Peng Cheng Laboratory*; Zibin Zheng, *Sun Yat-sen University*

PPipe: Efficient Video Analytics Serving on Heterogeneous GPU Clusters via Pool-Based Pipeline Parallelism 679
Z. Jonny Kong, Qiang Xu, and Y. Charlie Hu, *Purdue University*

Voltrix: Sparse Matrix-Matrix Multiplication on Tensor Cores with Asynchronous and Balanced Kernel Optimization 699
Yaqi Xia and Weihua Wang, *Wuhan University*; Donglin Yang, *Nvidia Corporation*; Xiaobo Zhou, *University of Macau*;
Dazhao Cheng, *Wuhan University*

Resilient Systems: Failure Detection, Consistency, and Scalability

NetKeeper: Enhancing Network Resilience with Autonomous Network Configuration Update on Traffic Patterns and Anomalies 715
Zhaoyang Wan, Rongxin Han, Haifeng Sun, Qi Qi, Zirui Zhuang, and Bo He, *State Key Laboratory of Networking and Switching Technology, Beijing University of Posts and Telecommunications*; Liang Zhang, *Huawei Technologies Co., Ltd*; Jianxin Liao and Jingyu Wang, *State Key Laboratory of Networking and Switching Technology, Beijing University of Posts and Telecommunications*

GREYHOUND: Hunting Fail-Slows in Hybrid-Parallel Training at Scale 731
Tianyuan Wu and Wei Wang, *Hong Kong University of Science and Technology*; Yinghao Yu, Siran Yang, and Wenchao Wu, *Alibaba Group*; Qinkai Duan, *Hong Kong University of Science and Technology*; Guodong Yang, Jiamang Wang, Lin Qu, and Liping Zhang, *Alibaba Group*

Crash Consistency in Block-Level Caching Systems: An Open CAS Case Study 749
Shaohua Duan, *Washington State University*; Youmin Chen, *Shanghai Jiao Tong University*

FiDe: Reliable and Fast Crash Failure Detection to Boost Datacenter Coordination 765
Davide Rovelli, *Università della Svizzera Italiana and SAP SE*; Pavel Chuprikov, *Télécom Paris and Institut Polytechnique de Paris*; Philipp Berdesinski, *turba*; Ali Pahlevan, *SAP SE*; Patrick Jahnke, *turba*; Patrick Eugster, *Università della Svizzera Italiana*

Network Performance & Protocols: From Space to VR

LEOCraft: Towards Designing Performant LEO Networks 789
Suvam Basak and Amitangshu Pal, *Indian Institute of Technology Kanpur*; Debopam Bhattacharjee, *Microsoft Research India*

| | |
|---|-------------|
| Emulating Space Computing Networks with RHONE | 815 |
| <i>Liyang Wang, Peking University; Qing Li, Beijing University of Posts and Telecommunications; Yuhan Zhou and Zhaofeng Luo, Peking University; Donghao Zhang and Shangguang Wang, Beijing University of Posts and Telecommunications; Xuanzhe Liu and Chenren Xu, Peking University; and Key Laboratory of High Confidence Software Technologies, Ministry of Education (PKU)</i> | |
| Roaming Free in the VR World with MP² | 833 |
| <i>Yifei Xu, University of California, Los Angeles; Xumiao Zhang, University of Michigan and Alibaba Cloud; Yuning Chen, University of California, Merced; Pan Hu, Uber Technologies, Inc.; Xuan Zeng, Zhilong Zheng, Xianshang Lin, and Yanmei Liu, Alibaba Cloud; Songwu Lu, University of California, Los Angeles; Z. Morley Mao, University of Michigan; Wan Du, University of California, Merced; Dennis Cai and Ennan Zhai, Alibaba Cloud; Yunfei Ma, Uber Technologies, Inc.</i> | |
| STORM: a Multipath QUIC Scheduler for Quick Streaming Media Transport under Unstable Mobile Networks | 851 |
| <i>Liekun Hu and Changlong Li, East China Normal University, Jianghuai Advance Technology Center, and MoE Engineering Research Center of Hardware/Software Co-Design Technology and Application</i> | |
| Internet Connection Splitting: What's Old is New Again | 867 |
| <i>Gina Yuan, Thea Rossman, and Keith Winstein, Stanford University</i> | |
| Distributed Systems: Communication, Consensus, and Data Structures | |
| WIC: Hiding Producer-Consumer Synchronization Delays with Warp-Level Interrupt-based GPU Communications | 889 |
| <i>Jiajian Zhang, Xi'an Jiaotong-Liverpool University and University of Liverpool; Fangyu Wu, Xi'an Jiaotong-Liverpool University; Hai Jiang, Beijing University of Posts and Telecommunications; Qiufeng Wang, Xi'an Jiaotong-Liverpool University; Genlang Chen and Chaoyi Pang, NingboTech University</i> | |
| Primus: Unified Training System for Large-Scale Deep Learning Recommendation Models | 905 |
| <i>Jixi Shan, ByteDance Inc.; Xiuqi Huang, Zhejiang University; Yang Guo, Hongyue Mao, Ho-Pang Hsu, Hang Cheng, Can Wang, and Jun Song, ByteDance Inc.; Rui Shi, Bytedance Inc.; Xiaofeng Gao, Shanghai Jiao Tong University; Jingwei Xu, Shiru Ren, Jiaxiao Zheng, Hua Huang, Lele Yu, and Peng Xu, ByteDance Inc.; Guihai Chen, Shanghai Jiao Tong University</i> | |
| Chitu: Avoiding Unnecessary Fallback in Byzantine Consensus | 923 |
| <i>Rongji Huang, Xiangzhe Wang, Xiaofeng Yan, and Lei Fan, Shanghai Jiao Tong University; Guangtao Xue and Shengyun Liu, Shanghai Jiao Tong University and Shanghai Key Laboratory of Trusted Data Circulation, Governance and Web3</i> | |
| Fast Distributed Transactions for RDMA-based Disaggregated Memory | 943 |
| <i>Haodi Lu, Haikun Liu, Yujian Zhang, Zhuohui Duan, Xiaofei Liao, Hai Jin, and Yu Zhang, Huazhong University of Science and Technology</i> | |
| Cuckoo for Clients: Disaggregated Cuckoo Hashing | 959 |
| <i>Stewart Grant and Alex C. Snoeren, UC San Diego</i> | |
| Virtualization and Isolation: Security, Sharing, and Performance | |
| LITESHIELD: Secure Containers via Lightweight, Composable Userspace μKernel Services | 973 |
| <i>Kaesi Manakkal, The University of Texas at Arlington; Nathan Daughety and Marcus Pendleton, Air Force Research Laboratory (AFRL); Hui Lu, The University of Texas at Arlington</i> | |
| Accelerating Nested Virtualization with HyperTurtle | 987 |
| <i>Ori Ben Zur and Jakob Krebs, Technion - Israel Institute of Technology; Shai Aviram Bergman, Huawei Zurich Research Center; Mark Silberstein, Technion - Israel Institute of Technology</i> | |
| Efficient Performance-Aware GPU Sharing with Compatibility and Isolation through Kernel Space Interception | 1003 |
| <i>Shulai Zhang, Ao Xu, Quan Chen, Han Zhao, and Weihao Cui, Shanghai Jiao Tong University; Zhen Wang, Yan Li, and Limin Xiao, Lenovo; Minyi Guo, Shanghai Jiao Tong University</i> | |

AnchorNet: Bridging Live and Collaborative Streaming with a Unified Architecture1021
Tong Meng, Wei Zhang, Dong Chen, Zhen Wang, Quanqing Li, Changqing Yan, Wei Yang, Chao Yuan, Le Zhang, Jianxin Kuang, and Jianlin Xu, *ByteDance*

Scaling Complex Models: Distribution, Heterogeneity, and Efficiency

KATZ: Efficient Workflow Serving for Diffusion Models with Many Adapters 1037
Suyi Li, Lingyun Yang, Xiaoxiao Jiang, Hanfeng Lu, and Dakai An, *Hong Kong University of Science and Technology*; Zhipeng Di, Weiyi Lu, Jiawei Chen, Kan Liu, Yinghao Yu, Tao Lan, Guodong Yang, Lin Qu, and Liping Zhang, *Alibaba Group*; Wei Wang, *Hong Kong University of Science and Technology*

PopFetcher: Towards Accelerated Mixture-of-Experts Training Via Popularity Based Expert-Wise Prefetch 1053
Junyi Zhang, Chuanhu Ma, Xiong Wang, and Yuntao Nie, *Huazhong University of Science and Technology*; Yuqing Li, *Wuhan University*; Yuedong Xu, *Fudan University*; Xiaofei Liao, *Huazhong University of Science and Technology*; Bo Li, *Hong Kong University of Science and Technology*; Hai Jin, *Huazhong University of Science and Technology*

HYPRECA: Distributed Heterogeneous In-Memory Embedding Database for Training Recommender Models 1071
Jiaao He, Shengqi Chen, Kezhao Huang, and Jidong Zhai, *Tsinghua University*

CrossPipe: Towards Optimal Pipeline Schedules for Cross-Datacenter Training 1089
Tiancheng Chen, Ales Kubicek, Langwen Huang, and Torsten Hoefler, *ETH Zurich*

Wednesday, July 9

Hunting Elusive Bugs: Verification and Analysis from Compilers to Hardware

Unveiling Compiler Faults via Attribute-Guided Compilation Space Exploration 1109
Jiangchang Wu, Yibiao Yang, Maolin Sun, and Yuming Zhou, *State Key Laboratory for Novel Software Technology, Nanjing University*

Understanding and Detecting Fail-Slow Hardware Failure Bugs in Cloud Systems 1127
Gen Dong and Yu Hua, *Huazhong University of Science and Technology*; Yongle Zhang, *Purdue University*; Zhangyu Chen and Menglei Chen, *Huazhong University of Science and Technology*

CONVEROS: Practical Model Checking for Verifying Rust OS Kernel Concurrency1143
Ruize Tang, *State Key Laboratory for Novel Software Technology, Nanjing University*; Minghua Wang, *Ant Group*; Xudong Sun, *University of Illinois Urbana-Champaign*; Lin Huang, *Ant Group*; Yu Huang and Xiaoxing Ma, *State Key Laboratory for Novel Software Technology, Nanjing University*

BIN2WRONG: a Unified Fuzzing Framework for Uncovering Semantic Errors in Binary-to-C Decompilers1161
Zao Yang and Stefan Nagy, *University of Utah*

HEC: Equivalence Verification Checking for Code Transformation via Equality Saturation1181
Jiaqi Yin and Zhan Song, *University of Maryland, College Park*; Nicolas Bohm Agostini and Antonino Tumeo, *Pacific Northwest National Laboratory*; Cunxi Yu, *University of Maryland, College Park*

Software-Hardware Synergy: Accelerators, Memory, and Interconnects

Para-ksm: Parallelized Memory Deduplication with Data Streaming Accelerator 1197
Houxiang Ji, *University of Illinois Urbana-Champaign*; Minh Kim and Seonmu Oh, *Daegu Gyeongbuk Institute of Science and Technology*; Daehoon Kim, *Yonsei University*; Nam Sung Kim, *University of Illinois Urbana-Champaign*

DSA-2LM: A CPU-Free Tiered Memory Architecture with Intel DSA 1213
Ruili Liu, *Tsinghua University and University of Electronic Science and Technology of China*; Teng Ma, *Alibaba Group*; Mingxing Zhang, Jialiang Huang, and Yingdi Shan, *Tsinghua University*; Zheng Liu, *Alibaba Group*; Lingfeng Xiang, Zhen Lin, Hui Lu, and Jia Rao, *The University of Texas at Arlington*; Kang Chen and Yongwei Wu, *Tsinghua University*

Turbocharge ANNS on Real Processing-in-Memory by Enabling Fine-Grained Per-PIM-Core Scheduling 1223
Puqing Wu, Minhui Xie, Enrui Zhao, Dafang Zhang, and Jing Wang, *Renmin University of China*; Xiao Liang and Kai Ren, *Kuaishou*; Yunpeng Chai, *Renmin University of China*

| | |
|---|-------------|
| SwCC: Software-Programmable and Per-Packet Congestion Control in RDMA Engine | 1243 |
| Hongjing Huang, Jie Zhang, Xuzheng Chen, Ziyu Song, Jiajun Qin, and Zeke Wang, <i>Zhejiang University</i> | |
| DRack: A CXL-Disaggregated Rack Architecture to Boost Inter-Rack Communication | 1261 |
| Xu Zhang and Ke Liu, <i>SKLP, Institute of Computing Technology, CAS; and University of Chinese Academy of Sciences</i> ; Yuan Hui and Xiaolong Zheng, <i>Huawei</i> ; Yisong Chang, <i>SKLP, Institute of Computing Technology, CAS; and University of Chinese Academy of Sciences</i> ; Yizhou Shan, <i>Huawei Cloud</i> ; Guanghui Zhang, <i>Shandong University</i> ; Ke Zhang, Yungang Bao, Mingyu Chen, and Chenxi Wang, <i>SKLP, Institute of Computing Technology, CAS; and University of Chinese Academy of Sciences</i> | |
| Securing the Stack: Attestation, Memory Protection, and Privacy | |
| ShieldReduce: Fine-Grained Shielded Data Reduction | 1281 |
| Jingyuan Yang, Jun Wu, Ruilin Wu, and Jingwei Li, <i>University of Electronic Science and Technology of China</i> ; Patrick P. C. Lee, <i>The Chinese University of Hong Kong</i> ; Xiong Li and Xiaosong Zhang, <i>University of Electronic Science and Technology of China</i> | |
| MemoryTrap: Booby Trapping Memory to Counter Memory Disclosure Attacks with Hardware Support | 1297 |
| Chenke Luo, <i>Wuhan University and Tulane University</i> ; Jiang Ming, <i>Tulane University</i> ; Dongpeng Xu, <i>University of New Hampshire</i> ; Guojun Peng and Jianming Fu, <i>Wuhan University</i> | |
| Separate but Together: Integrating Remote Attestation into TLS | 1319 |
| Carsten Weinhold, <i>Barkhausen Institut</i> ; Muhammad Usama Sardar, <i>TU Dresden</i> ; Ionut Mihalcea and Yogesh Deshpande, <i>Arm</i> ; Hannes Tschofenig, <i>University of Applied Sciences Bonn-Rhein-Sieg</i> ; Yaron Sheffer, <i>Intuit</i> ; Thomas Fossati, <i>Linaro</i> ; Michael Roitzsch, <i>Barkhausen Institut</i> | |
| DDLUMOS: Understanding and Detecting Atomic DDL Bugs in DBMSs | 1327 |
| Zhiyong Wu, <i>Tsinghua University</i> ; Jie Liang, <i>Beihang University</i> ; Jingzhou Fu, Wenqian Deng, and Yu Jiang, <i>Tsinghua University</i> | |
| Hardware-Specific Optimizations: Space, Accelerators, and AI Chips | |
| SpaceExit: Enabling Efficient Adaptive Computing in Space with Early Exits | 1343 |
| Jiacheng Liu, <i>Shanghai Jiao Tong University and The Chinese University of Hong Kong</i> ; Xiaozhi Zhu, Tongqiao Xu, Xiaofeng Hou, and Chao Li, <i>Shanghai Jiao Tong University</i> | |
| XRT: An Accelerator-Aware Runtime for Accelerated Chip Multiprocessors | 1359 |
| Neel Patel and Mohammad Alian, <i>Cornell University</i> | |
| DShuffle: DPU-Optimized Shuffle Framework for Large-scale Data Processing | 1371 |
| Chen Ding, Sicen Li, and Kai Lu, <i>Wuhan National Laboratory for Optoelectronics, Huazhong University of Science and Technology</i> ; Ting Yao, Daohui Wang, and Huatao Wu, <i>Huawei Cloud</i> ; Jiguang Wan, Zhihu Tan, Changsheng Xie, <i>Wuhan National Laboratory for Optoelectronics, Huazhong University of Science and Technology</i> | |
| Accelerating Model Training on Ascend Chips: An Industrial System for Profiling, Analysis and Optimization | 1387 |
| Yuhang Zhou, Zibo Wang, Zhibin Wang, Ruyi Zhang, Chen Tian, Xiaoliang Wang, Wanchun Dou, and Guihai Chen, <i>Nanjing University</i> ; Bingqiang Wang, Yonghong Tian, Yan Zhang, and Hui Wang, <i>Peng Cheng Laboratory</i> ; Fuchun Wei, Boquan Sun, Jingyi Zhang, Bin She, Teng Su, Yifan Yao, Chunsheng Li, Ziyang Zhang, and Yaoyuan Wang, <i>Huawei</i> ; Bin Zhou, <i>Shandong University</i> ; Guyue Liu, <i>Peking University</i> | |
| Finding Faults: Concurrency, Numerics, and Kernel Fuzzing | |
| CAFault: Enhance Fault Injection Technique in Practical Distributed Systems via Abundant Fault-Dependent Configurations | 1409 |
| Yuanliang Chen, Fuchen Ma, Yuanhang Zhou, Zhen Yan, and Yu Jiang, <i>Tsinghua University</i> | |
| Revealing Floating-Point Accumulation Orders in Software/Hardware Implementations | 1425 |
| Peichen Xie, Yanjie Gao, Yang Wang, and Jilong Xue, <i>Microsoft Research</i> | |
| Inferring Likely Counting-related Atomicity Program Properties for Persistent Memory | 1441 |
| Yunmo Zhang and Junqiao Qiu, <i>City University of Hong Kong</i> ; Hong Xu, <i>The Chinese University of Hong Kong</i> ; Chun Jason Xue, <i>MBZUAI</i> | |

Optimizing Input Minimization in Kernel Fuzzing1451
Hui Guo, *East China Normal University*; Hao Sun, *ETH Zurich*; Shan Huang, Ting Su, and Geguang Pu, *East China Normal University*; Shaohua Li, *The Chinese University of Hong Kong*

IRHash: Efficient Multi-Language Compiler Caching by IR-Level Hashing 1467
Tobias Landsberg and Johannes Grunenber, *Leibniz Universität Hannover*; Christian Dietrich, *Technische Universität Braunschweig*; Daniel Lohmann, *Leibniz Universität Hannover*

Advanced Distributed Systems: ML, Data, and Performance Tuning

On-Demand Container Partitioning for Distributed ML..... 1481
Giovanni Bartolomeo, Navidreza Asadi, Wolfgang Kellerer, and Jorg Ott, *Technical University of Munich*; Nitinder Mohan, *TU Delft*

PathWeaver: A High-Throughput Multi-GPU System for Graph-Based Approximate Nearest Neighbor Search 1501
Sukjin Kim, Seongyeon Park, Si Ung Noh, Junguk Hong, Taehee Kwon, Hunseong Lim, and Jinho Lee, *Seoul National University*

Universal Checkpointing: A Flexible and Efficient Distributed Checkpointing System for Large-Scale DNN Training with Reconfigurable Parallelism..... 1519
Xinyu Lian, *University of Illinois Urbana–Champaign*; Sam Ade Jacobs, Lev Kurilenko, and Masahiro Tanaka, *Microsoft*; Stas Bekman, *Snowflake*; Olatunji Ruwase, *Microsoft*; Minjia Zhang, *University of Illinois Urbana–Champaign*

Towards High-Performance Transactional Stateful Serverless Workflows with Affinity-Aware Leasing 1535
Jianjun Zhao, Haikun Liu, Shuhao Zhang, and Haodi Lu, *Huazhong University of Science and Technology*; Yancan Mao, *School of Computing, National University of Singapore*; Zhuohui Duan, Xiaofei Liao, and Hai Jin, *Huazhong University of Science and Technology*

Swift: Fast Performance Tuning with GAN-Generated Configurations..... 1553
Chao Chen and Shixin Huang, *Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences*; Xuehai Qian, *Tsinghua University*; Zhibin Yu, *Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences*; and Shuhai Lab, *Huawei Cloud*

OS, Mobile, and Reliability Challenges

PMR: Fast Application Response via Parallel Memory Reclaim on Mobile Devices 1569
Wentong Li, *MoE Engineering Research Center of Software/Hardware Co-Design Technology and Application, East China Normal University, Shanghai, China*; and *School of Computer Science and Technology, East China Normal University*; Li-Pin Chang, *National Yang Ming Chiao Tung University, Taiwan*; Yu Mao, *City University of Hong Kong, Hong Kong*; and *MBZUAI, Abu Dhabi*; Liang Shi, *MoE Engineering Research Center of Software/Hardware Co-Design Technology and Application, East China Normal University, Shanghai, China*; and *School of Computer Science and Technology, East China Normal University*

SAVE: Software-Implemented Fault Tolerance for Model Inference against GPU Memory Bit Flips 1585
Wenxin Zheng, Bin Xu, Jinyu Gu, and Haibo Chen, *Shanghai Jiao Tong University*

Identifying and Analyzing Pitfalls in GNN Systems 1605
Yidong Gong, Arnab Kanti Tarafder, Saima Afrin, and Pradeep Kumar, *William & Mary*

Bluetooth Low Energy Security Testing with Combinatorial Methods1625
Dominik-Philip Schreiber, Manuel Leithner, and Jovan Zivanovic, *SBA Research*; Dimitris E. Simos, *SBA Research, Salzburg University of Applied Sciences, and Paris Lodron University of Salzburg*

Intelligent Resource Management: Federation, Colocation, and ML for Systems

Resource Multiplexing in Tuning and Serving Large Language Models 1639
Yongjun He and Haofeng Yang, *ETH Zurich*; Yao Lu, *National University of Singapore*; Ana Klimovic and Gustavo Alonso, *ETH Zurich*

Colocating ML Inference and Training with Fast GPU Memory Handover 1657
Jiali Wang, Yankui Wang, Mingcong Han, and Rong Chen, *Institute of Parallel and Distributed Systems, Shanghai Jiao Tong University*

| | |
|---|-------------|
| AssyLLM: Efficient Federated Fine-tuning of LLMs via Assembling Pre-trained Blocks | 1677 |
| Shichen Zhan, Li Li, and Chengzhong Xu, <i>University of Macau</i> | |
| Learning-Enhanced High-Throughput Pattern Matching Based on Programmable Data Plane | 1693 |
| Guanglin Duan and Yucheng Huang, <i>Peng Cheng Laboratory, Tsinghua Shenzhen International Graduate School</i> ; Zhengxin Zhang, <i>Peng Cheng Laboratory, Tsinghua Shenzhen International Graduate School, and Cornell University</i> ; Qing Li and Dan Zhao, <i>Peng Cheng Laboratory</i> ; Zili Meng, <i>Hong Kong University of Science and Technology</i> ; Dirk Kutscher, <i>Hong Kong University of Science and Technology (Guangzhou)</i> ; Ruoyu Li, <i>Shenzhen University and Peng Cheng Laboratory</i> ; Yong Jiang, <i>Tsinghua Shenzhen International Graduate School</i> ; Mingwei Xu, <i>Tsinghua University</i> | |