

2024 International Conference on Networking, Architecture and Storage (NAS 2024)

**Zhuhai, China
9-11 November 2024**



**IEEE Catalog Number: CFP2462C-POD
ISBN: 979-8-3315-2051-9**

**Copyright © 2024 by the Institute of Electrical and Electronics Engineers, Inc.
All Rights Reserved**

Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limit of U.S. copyright law for private use of patrons those articles in this volume that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

For other copying, reprint or republication permission, write to IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08854. All rights reserved.

****** This is a print representation of what appears in the IEEE Digital Library. Some format issues inherent in the e-media version may also appear in this print version.***

IEEE Catalog Number:	CFP2462C-POD
ISBN (Print-On-Demand):	979-8-3315-2051-9
ISBN (Online):	979-8-3315-2050-2
ISSN:	2835-334X

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
E-mail: curran@proceedings.com
Web: www.proceedings.com

CURRAN ASSOCIATES INC.
proceedings
.com

Table of Contents

Machine learning and Accelerators

DeepFetch: A Node-Aware Greedy Fetch System for Distributed Cache of Deep Learning Applications	1
<i>Lijuan Kong, Fei Mei, Chunjie Zhu, Wen Cheng, and Lingfang Zeng</i>	
IOWA: An I/O-Aware Adaptive Sampling Framework for Deep Learning.....	9
<i>Shuang Hu, Weijian Chen, Yanlong Yin, and Shuibing He</i>	
mLOOP: Optimize Loop Unrolling in Compilation with a ML-Based Approach	17
<i>Zhongchun Zheng, Yuan Wu, and Xianwei Zhang</i>	
Accurate Generation of I/O Workloads using Generative Adversarial Networks	25
<i>Heyu Zhang, Zhen Yang, Yulai Xie, Yafeng Wu, Jiakun Li, Dan Feng, Avani Wildani, and Darrell Long</i>	
Hardware and Software Co-Design of FPGA-Based Erasure Coding Accelerator for Enterprise Storage.....	34
<i>Xueqing Zhang, Xianwei Meng, Tao Liu, Bole Li, Bin Sun, and Hui Li</i>	
SAQO: Empowering Computational Storage Device for Efficient SQL Query Acceleration	40
<i>Yao Deng, Pengze Lv, Mengran Zhang, and Wei Tong</i>	
RL-Based Scheduling and Placement for Deep Learning Jobs on Large-Scale GPU Clusters.....	44
<i>Jiayuan Liao and Minghua Shen</i>	
Network and Storage	
EWS: Towards Cost-Effective Job Scheduling via Combinatorial Multi-Armed Bandit Learning	48
<i>Linchang Xiao, Zili Xiao, Di Wu, and Miao Hu</i>	
CGHit: A Content-Oriented Generative-Hit Framework for Content Delivery Networks	56
<i>Peng Wang, Yu Liu, Kai Han, Ziqi Liu, Ke Liu, Mingyang Wang, Ke Zhou, and Zhihai Huang</i>	
TranLogs: Lossless Failure Recovery Empowered by Training Logs	64
<i>Xiaoyu Liu and Lingfang Zeng</i>	
OCSL: An Online Compression Scheme for Streaming Semi-Structured Logs.....	72
<i>Qing Xie, Zhiye Li, Meiqiong Yuan, Chunhua Li, and Ke Zhou</i>	
MPVSched: Multipath Transmissions and Video Frame Scheduling for Content Delivery Networks.....	80
<i>Haoran Xu, Xiaoxi Zhang, Jingpu Duan, Chuan Wu, Jinhang Zuo, Xuan Zeng, Chen Chen, Yubing Qiu, and Xu Chen</i>	
ERD: AVX-512-Based Enhancement of Resemblance Detection for Post-Deduplication Delta Compression	84
<i>Tianyun Chen, Hao Liu, Fang Wang, Weijie Yan, Ningyuan Mo, and Xiaolei Liu</i>	

F4: Fast, Flexible and Stable-Fortified User-Space Thread Scheduling Framework..... 88
Xiaobo Zheng, Zihao Duan, Shiyi Li, Haojun Hu, and Wen Xia

Distributed Systems

Erasur Coding Based Optimization in Decentralized Distributed Storage Systems..... 96
Yiwei Gan, Zhijie Huang, Yulong Shi, Xiao Zhang, and Nannan Zhao

A New Family of Wide Locally Repairable Codes for Seeking the Trade-off between Fault Tolerance and Repair Degree in Distributed Storage System 104
Yan Wang, Junhao Shi, Hao Wang, and Yanghuang Cao

POFFO: A Perceptual Online File Fingerprint Offloading Strategy for Effective Data Deduplication at Cloud-Edge Systems..... 112
Hexian Lu, Yuhui Deng, and Jiande Huang

Design and Implementation of a Turbulence Data Sharing Platform for Scientific Big Data 120
Youjun Zhao, Xiao Zhang, Wendi Cheng, Zhaohui Pan, Chenguang Sun, and Xueqiang Shan

A Unified Computational Storage and Memory Architecture in the CXL Era 128
Xiang Chen, Fei Wu, and Tao Lu

Composable Storage Servers: A Storage Paradigm for Disaggregated Systems..... 132
Shai Bergman, Onur Mutlu, Wu Yong, Keji Huang, and Ji Zhang

CAMS: A Cost-Aware Migration Scheme for Cloud Object Storage Systems 136
Man Luo, Ke Zhou, Hua Wang, Zhibin Jiang, Han Kong, and Yongguang Ji

GooseFS: Distributed Cache Service to Enhance Cloud Object Storage Performance 140
Zheng Chen, Li Cheng, Junming Yan, Xuefeng Li, Zhitao Chen, Wei Zhang, Honggang Chen, Yang Yu, Zheng Han, and Zhiwei Yang

Solid State Drives Architecture and Systems

SSD Failures in Large-Scale Data Centers: What? Why? and How? 144
Wenyan You, Jiayuan Dong, Xingdi Feng, Zeyun Chen, Bo Mao, and Suzhen Wu

Improving F2FS fsync() Latency through Parallelizing Dnode and Data Page Writeback..... 152
Mengyang Ma, Yumiao Zhao, Yunpeng Song, Yi Zhang, and Shouzhen Gu

Zoned-WB: WriteBooster Design with Zoned Storage for User Experience on Smartphones 160
Dingcui Yu, Ziang Huang, Wentong Li, Zonghuan Yan, Shouzhen Gu, and Liang Shi

In-Place Switch: Reprogramming Based SLC Cache Design for Hybrid 3D SSDs..... 164
Xufeng Yang, Jiancong Zheng, Cheng Ji, and Congming Gao

MCBGC: A Multi-Threshold Copyback-Based Garbage Collection Scheme for 3D NAND Flash Memory..... 168
Kang Hu, Haihua Hu, Pengchao Han, and Guojun Han

LA-Write: Balancing Endurance of Inter-Layer for Prolonging 3D NAND Flash Memory Lifetime 172
Yu Wan, Yajuan Du, and Siyi Huang

Dynamic Cache Partitioning for Enhancing Parallel I/O Performance in NVMe SSDs	176
<i>Zecheng Li, Guancheng Li, Songhui Cao, Shu Yin, and Xiaojun Ruan</i>	
Author Index.....	180