

# **2021 IEEE International Conference on Networking, Architecture and Storage (NAS 2021)**

**Riverside, California, USA  
24-26 October 2021**



**IEEE Catalog Number: CFP2162C-POD  
ISBN: 978-1-7281-7745-8**

**Copyright © 2021 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:	CFP2162C-POD
ISBN (Print-On-Demand):	978-1-7281-7745-8
ISBN (Online):	978-1-7281-7744-1

**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](mailto:curran@proceedings.com)  
Web: [www.proceedings.com](http://www.proceedings.com)

CURRAN ASSOCIATES INC.  
**proceedings**  
.com

# Table of Contents

Message from the General Chair .....	vii
Message from Program Chairs .....	viii
Sponsors .....	ix

## Industry Session

Machine Reasoning – Improving AIOps for Intent Based Networks..... <i>Pradeep Kathail, CTO Cisco</i>	x
World’s #1 CRM Scale Challenges..... <i>Jason Ding, SVP Engineering, Salesforce</i>	xi
Making Enterprise SSD Smarter and Faster..... <i>Yafei Yang CEO, Shenzhen Dapu Microelectronics Co., Ltd</i>	xii

## Session 1: Architecture

Cache Compression with Efficient in-SRAM Data Comparison..... <i>Xiaowei Wang, Charles Augustine, Eriko Nurvitadhi, Li Zhao, Ravi Iyer, and Reetuparna Das</i>	1
GO: Out-of-Core Partitioning of Large Irregular Graphs..... <i>Rajiv Gupta and Gurneet Kaur</i>	9
PowerPrep: A Power Management Proposal for User-Facing Datacenter Workloads .....	19
<i>Vineetha Govindaraj, Sumitha George, Mahmut Kandemir, John Sampson, and Vijaykrishnan Naryanan</i>	
Efficient NVM Crash Consistency by Mitigating Resource Contention..... <i>Zhiyuan Lu, Jianhui Yue, Yifu Deng, and Yifeng Zhu</i>	26
Migrating Software from x86 to ARM Architecture: An Instruction Prediction Approach..... <i>Blake W. Ford, Apan Qasem, Jelena Tesic, and Ziliang Zong</i>	34

## Session 2: Storage

A Comprehensive Empirical Study of File Systems on Optane Persistent Memory..... <i>Yang Yang, Qiang Cao, and Shengjue Wang</i>	40
NVSwap: Latency-Aware Paging using Non-Volatile Main Memory..... <i>Yekang Wu and Xuechen Zhang</i>	48
Constant Time Garbage Collection in SSDs .....	57
<i>Reza Salkhordeh, Kevin Kremer, Lars Nagel, Dennis Maisenbacher, Hans Holmberg, Matias Bjørling, and Andre Brinkmann</i>	

Edges: Evenly Distributing Garbage-Collections for Enterprise SSDs via Stochastic Optimization ..... 66  
*Shuyi Pei, Jing Yang, and Bin Li*

A New PIS Accelerator for Text Searching ..... 70  
*Yunxin Huang, Aiguo Song, and Yafei Yang*

### **Session 3: Network**

Minimizing the Number of Rules to Mitigate Link Congestion in SDN-Based Datacenters..... 76  
*Rajorshi Biswas and Jie Wu*

Design of A Multi-Path Reconfigurable Traffic Monitoring System..... 84  
*Liang-Min Wang, Tim Miskell, John Morgan, and Edwin Verplanke*

Terrestrial and Space-Based Cloud Computing with Scalable, Responsible and Explainable Artificial Intelligence - A Position Paper ..... 92  
*Davide Martizzi and Partha Datta Ray*

Implementation of a High-Throughput Virtual Switch Port Monitoring System..... 96  
*Tim Miskell, Patrick Fu, Liang-Min Wang, Cunming Liang, and Edwin Verplanke*

An Incast-Coflow-Aware Minimum-Rate-Guaranteed Congestion Control Protocol for Datacenter Applications ..... 104  
*Zhijun Wang, Yunxiang Wu, Rosenkrantz Stoddard, Ning Li, Minh Nguyen, and Hao Che*

Balancing Latency and Quality in Web Search..... 112  
*Liang Zhou and K. K. Ramakrishnan*

### **Session 4: Architecture**

Vulkan vs OpenGL ES: Performance and Energy Efficiency Comparison on the big.LITTLE Architecture ..... 116  
*Michael Lujan, Michael McCrary, Blake W. Ford, and Ziliang Zong*

LocalityGuru: A PTX Analyzer for Extracting Thread Block-Level Locality in GPGPUs ..... 124  
*Devashree Tripathy, Amir Ali Abdolrashidi, Quan Fan, Daniel Wong, and Manoranjan Satpathy*

GPU-Assisted Memory Expansion..... 132  
*Pisacha Srinuan, Purushottam Sigdel, Xu Yuan, Lu Peng, Paul Darby, Christopher Aucoin, and Nian-Feng Tzeng*

E2E Visual Analytics: Achieving >10X Edge/Cloud Optimizations ..... 140  
*Chaunte Lacewell, Nilesh Ahuja, Pablo Munoz, Parual Datta, Ragaad Altarawneh, Vui Seng Chua, Nilesh Jain, Omesh Tickoo, and Ravi Iyer*

Locality-Aware Thread Block Design in Single and Multi-GPU Graph Processing ..... 148  
*Quan Fan and Zizhong Chen*

MLPP: Exploring Transfer Learning and Model Distillation for Predicting Application Performance ..... 152  
*Jashwant Raj Gunasekaran and Cyan Subhra Mishra*

## Session 5: Architecture

ICAP: Designing Inrush Current Aware Power Gating Switch for GPGPU.....	156
<i>Hadi Zamani, Devashree Tripathy, Ali Jahanshahi, and Daniel Wong</i>	
Machine Learning-Based Vulnerability Study of Interpose PUFs as Security Primitives for IoT Networks .....	164
<i>Bipana Thapaliya, Khalid T. Mursi, and Yu Zhuang</i>	
Reducing the Training Overhead of the HPC Compression Autoencoder via Dataset Proportioning .....	171
<i>Tong Liu, Shakeel Alibhai, Jinzhen Wang, Qing Liu, and Xubin He</i>	
Exploring Storage Device Characteristics of a RISC-V Little-core SoC.....	178
<i>Tao Lu</i>	
Deflection-Aware Routing Algorithm in Network on Chip against Soft Errors and Crosstalk Faults.....	182
<i>Hadi Zamani, Zahra Shirmohammadi, and Ali Jahanshahi</i>	

## Session 6: Storage

CALC: A Content-Aware Learning Cache for Storage Systems .....	188
<i>Maher Kachmar and David Kaeli</i>	
Improving Relational Database upon the Arrival of Storage Hardware with Built-in Transparent Compression.....	196
<i>Yifan Qiao, Xubin Chen, Jingpeng Hao, Jiangpeng Li, Jingqiang Wang, Qi Wu, Yang Liu, and Tong Zhang</i>	
On Adapting the Cache Block Size in SSD Caches .....	205
<i>Nikolaus Jeremic, Helge Parzyjegl, and Gero Muhl</i>	
Cached Mapping Table Prefetching for Random Reads in Solid-State Drives.....	213
<i>Xiaojun Ruan, Xunfei Jiang, and Haiquan Chen</i>	
EFLOG: A Full Stream-Logging Scheme with Erasure Coding in Cloud Storage Systems.....	219
<i>Lei Sun, Qiang Cao, Shucheng Wang, and Changsheng Xie</i>	

## Session 7: Network

Decoupling Control and Data Transmission in RDMA Enabled Cloud Data Centers .....	223
<i>Peter Varman and Qingyue Liu</i>	
Leveraging Network Delay Variability to Improve QoE of Latency Critical Services.....	231
<i>Matthew K. Farrens and Sambit K. Shukla</i>	
Egress Peer Engineering over MPLS Label Unicast Transport.....	239
<i>Sundaram Tirunelveli Radhakrishnan and Satya R. Mohanty</i>	
A Congestion Aware Multi-Path Label Switching in Data Centers using Programmable Switches.....	243
<i>Yeim-Kuan Chang, Hung-Yen Wang, and Yu-Hsiang Lin</i>	

Towards a Proactive Lightweight Serverless Edge Cloud for Internet-of-Things Applications .....	251
<i>Ian-Chin Wang, Shixiong Qi, Elizabeth Liri, and K. K. Ramakrishnan</i>	
Towards Energy-Efficient and Real-Time Cloud Computing.....	255
<i>Taha Tekreeti, Ting Cao, Xiaopu Peng, Tathagata Bhattacharya, Jianzhou Mao, Xiao Qin, and Wei-Shinn Ku</i>	
<b>Session 8: Storage</b>	
Flow Scheduling in a Heterogeneous NFV Environment using Reinforcement Learning.....	259
<i>Chun-Jen Lin, Yan Luo, Liang-Min Wang, and Li-De Chen</i>	
Implementing Flash-Cached Storage Systems using Computational Storage Drive with Built-in Transparent Compression.....	267
<i>Jingpeng Hao, Xubin Chen, Yifan Qiao, Yuyang Zhang, and Tong Zhang</i>	
Fast Variable-Grained Resemblance Data Deduplication for Cloud Storage.....	275
<i>Xuming Ye, Jia Tang, Wenlong Tian, Ruixuan Li, Weijun Xiao, Yuqing Geng, and Zhiyong Xu</i>	
A Case Study of Migrating RocksDB on Intel Optane Persistent Memory .....	283
<i>Ziyi Lu and Qiang Cao</i>	
PLMC: A Predictable Tail Latency Mode Coordinator for Shared NVMe SSD with Multiple Hosts .....	287
<i>Tanaya Roy, Jit Gupta, Krishna Kant, Amitangshu Pal, Dave Minturn, and Arash Tavakkol</i>	
Characterizing AI Model Inference Applications Running in the SGX Environment .....	293
<i>Shixiong Jing, Qinkun Bao, Pei Wang, Xulong Tang, and Dinghao Wu</i>	
Author Index.....	297