

# **2017 International Conference on Networking, Architecture, and Storage (NAS 2017)**

**Shenzhen, China  
7-9 August 2017**



**IEEE Catalog Number: CFP1762C-POD  
ISBN: 978-1-5386-3487-5**

**Copyright © 2017 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:	CFP1762C-POD
ISBN (Print-On-Demand):	978-1-5386-3487-5
ISBN (Online):	978-1-5386-3486-8

**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

<b>DD-L1D: Improving the Decoupled L1D Efficiency for GPU Architecture .....</b>	<b>1</b>
<i>Weiguang Yang, Yuxin Wang, Yulong Yu, Guangyuan Kan and He Guo</i>	
<b>TraceRAR: An I/O Performance Evaluation Tool for Replaying, Analyzing, and Regenerating Traces .....</b>	<b>11</b>
<i>Bingzhe Li, Farnaz Toussi, Clark Anderson, David J. Lilja, and David H.C. Du</i>	
<b>Towards Robust and Accurate Similar Trajectory Discovery: Weak-parametric Approaches.....</b>	<b>21</b>
<i>Yupeng Tuo, Xiaochun Yun, and Yongzheng Zhang</i>	
<b>A High-Performance Persistent Identifier Management Protocol .....</b>	<b>31</b>
<i>Fatih Berber and Ramin Yahyapour</i>	
<b>Performance Evaluation and Modeling of HPC I/O on Non-Volatile Memory .....</b>	<b>41</b>
<i>Wei Liu, Kai Wu, Jialin Liu, Feng Chen, and Dong Li</i>	
<b>WPS: A Workload-aware Placement Scheme for Erasure-Coded In-Memory Stores .....</b>	<b>51</b>
<i>Shuang Wang, Jianzhong Huang, Xiao Qin, Qiang Cao, and Changsheng Xie</i>	
<b>Reducing Chunk Fragmentation for In-line Delta Compressed and Deduplicated Backup Systems .....</b>	<b>61</b>
<i>Yucheng Zhang, Dan Feng, Yu Hua, Yuchong Hu, Wen Xia, Min Fu, Xiaolan Tang, Zhikun Wang, Fangting Huang, and Yukun Zhou</i>	
<b>Analysis and Correlation of Application I/O Performance and System-Wide I/O Activity .....</b>	<b>71</b>
<i>Sandeep Madireddy, Prasanna Balaprakash, Philip Carns, Robert Latham, Robert Ross, Shane Snyder, and Stefan M. Wild</i>	
<b>Detecting Information Theft Based on Mobile Network Flows for Android Users .....</b>	<b>81</b>
<i>Zhenyu Cheng, Xunxun Chen, Yongzheng Zhang, Shuhao Li, and Yafei Sang</i>	
<b>Service Migrations in the Cloud for Mobile Accesses: A Reinforcement Learning Approach.....</b>	<b>91</b>
<i>Shan Cao, Yang Wang, and Chengzhong Xu</i>	
<b>ISM - An Intra-Stripe Data Migration Approach for RAID-5 Scaling .....</b>	<b>101</b>
<i>Jie Liang, Yinlong Xu, Yongkun Li, and Yubiao Pan</i>	
<b>ALARM: A Location-Aware Redistribution Method to Improve 3D FG NAND Flash Reliability.....</b>	<b>111</b>
<i>Yue Zhu, Fei Wu, Qin Xiong, Zhonghai Lu, and Changsheng Xie</i>	
<b>Kaleido: Enabling Efficient Scientific Data Processing on Big-Data Systems.....</b>	<b>121</b>
<i>Saman Biookaghazadeh, Shujia Zhou, and Ming Zhao</i>	
<b>Megalloc*: Fast Distributed Memory Allocator for NVM-based Cluster .....</b>	<b>131</b>
<i>Songping Yu, Nong Xiao, Mingzhu Deng, Yuxuan Xing, Fang Liu, and Wei Chen</i>	
<b>Optimizing Energy Consumption on HPC Systems with a Multi-level Checkpointing Mechanism.....</b>	<b>140</b>
<i>Muhammad Alfian Amrizal and Hiroyuki Takizawa</i>	
<b>Contact Duration-Aware Routing in Delay Tolerant Networks.....</b>	<b>149</b>
<i>Tuan Le and Mario Gerla</i>	
<b>Data Block Partitioning for Recovering Stuck-at Faults in PCMs.....</b>	<b>157</b>
<i>Marjan Asadinia, Majid Jalili, and Hamid Sarbazi-Azad</i>	
<b>Extending Lifetime of SSD in RAID5 Systems through a Reliable Hierarchical Cache .....</b>	<b>165</b>
<i>Rui Ye, Wentao Meng, and Shenggang Wan</i>	

<b>A-MapCG: An Adaptive MapReduce Framework for GPUs.....</b>	<b>173</b>
<i>Lifeng Liu, Yue Zhang, Meilin Liu, Chongjun Wang, and Jun Wang</i>	
<b>MTM: A Reliable Multiple Trees Multicast for Data Center Network.....</b>	<b>181</b>
<i>Xin Xiong and Tan Chen</i>	
<b>Multicast Scheduling with Markov Chains in Fat-tree Data Center Networks.....</b>	<b>188</b>
<i>Guozhi Li, Songtao Guo, Guiyan Liu, and Yuanyuan Yang</i>	
<b>Binary Index and Journal Embedding in the Linear Tape File System .....</b>	<b>195</b>
<i>Klaus Birkelund Jensen and Brian Vinter</i>	
<b>Enhancing Next Generation Passive Optical Network Stage 2 (NG-PON2) with Channel Bonding .....</b>	<b>202</b>
<i>Liang Zhang, Yuanqiu Luo, Nirwan Ansari, Bo Gao, Xiang Liu, and Frank Effenberger</i>	
<b>A Write-Through Cache Method to Improve Small Write Performance of SSD-based RAID.....</b>	<b>208</b>
<i>Linjun Mei, Dan Feng, Jianxi Chen, Lingfang Zeng, and Jingning Liu</i>	
<b>Performance Tuning and Modeling for Big Data Applications in Docker Containers .....</b>	<b>214</b>
<i>Kejiang Ye and Yunjie Ji</i>	
<b>An Experimental Study on Deep Learning Based on Different Hardware Configurations.....</b>	<b>220</b>
<i>Jingjun Li, Chen Zhang, Qiang Cao, Chuanyi Qi, Jianzhong Huang, and Changsheng Xie</i>	
<b>Automatic Collector for Dynamic Cloud Performance Information .....</b>	<b>226</b>
<i>Olaf Elzinga, Spiros Koulouzis, Arie Taal, Junchao Wang, Yang Hu, Huan Zhou, Paul Martin, Cees de Laat, and Zhiming Zhao</i>	
<b>Facilitating Workload Aware Storage Platform by using Machine Learning Technics .....</b>	<b>232</b>
<i>Wubin Li, Fereydown Farrahi Moghaddam, Parisa Heidari, Yves Lemieux, and Abdelouahed Gherbi</i>	
<b>DualStack: A High Efficient Dynamic Page Scheduling Scheme in Hybrid Main Memory .....</b>	<b>238</b>
<i>Zhen Zhang, Yinjin Fu, and Guyu Hu</i>	
<b>A Hash-based Space-efficient Page-level FTL for Large-capacity SSDs .....</b>	<b>244</b>
<i>Fan Ni, Chunyi Liu, Yang Wang, Chengzhong Xu, Xiao Zhang, and Song Jiang</i>	
<b>Revisiting Updating Schemes for Erasure-coded In-Memory Stores .....</b>	<b>250</b>
<i>Jie Xia, Jianzhong Huang, Xiao Qin, Qiang Cao, and Changsheng Xie</i>	
<b>SpyStorage: A Highly Reliable Multi-Cloud Storage with Secure and Anonymous Data Sharing.....</b>	<b>256</b>
<i>Pengyan Shen, Wan Liu, Zheng Wu, Mingzhong Xiao, and Quanqing Xu</i>	
<b>Exploiting Virtual Metadata Servers to Provide Multi-level Consistency for Key-value Object-based Data Store .....</b>	<b>262</b>
<i>Xiaozhao Zhuang, Xiaoyang Qu, Zhiyong Lu, Jiguang Wan, and Changsheng Xie</i>	
<b>Balancing the Storage in a Deduplication Cluster .....</b>	<b>268</b>
<i>Giacomo Grangia, Quanqing Xu, Andrea Bianco, and Paolo Giaccone</i>	
<b>SimpleBP: A Lightweight Branch Prediction Simulator for Effective Design Exploration.....</b>	<b>272</b>
<i>Chaobing Zhou, Libo Huang, Zhisheng Li, and Qiang Dou</i>	
<b>hpFog: A FPGA-based Fog Computing Platform .....</b>	<b>274</b>
<i>Tze Hon Tan, Chia Yee Ooi, and M. N. Marsono</i>	
<b>Development of Network Simulator for LWA/LAA Implementations .....</b>	<b>276</b>
<i>Kun-Hee Lee, Min-Taek Choi, Sung-Chul Choi, and Jae-Hoon Kim</i>	

<b>Design of Universal Broker Architecture for Edge Networking .....</b>	<b>278</b>
<i>Joo-Hyun Kim, Seo-Hee Hong, So-Hyun Yang, and Jae-Hoon Kim</i>	
<b>Performance Optimization of In-Memory File System in Distributed Storage System .....</b>	<b>280</b>
<i>Zhaowei Li, Yunlong Yan, Jintao Mo, Zhaocong Wen, and Junmin Wu</i>	
<b>Branch Prediction Migration for Multi-core Architectures .....</b>	<b>282</b>
<i>Tan Zhang, Chaobing Zhou, Libo Huang, and Nong Xiao</i>	
<b>FlashStorageSim: Performance Modeling for SSD Architectures .....</b>	<b>284</b>
<i>Krishna T. Malladi, Mu-Tien Chang, Dimin Niu, and Hongzhong Zheng</i>	
<b>Rack Level Scheduling for Containerized Workloads .....</b>	<b>286</b>
<i>Qiumin Xu, Krishna T. Malladi, and Manu Awasthi</i>	
<b>Privacy-Protected Data Collection in Wireless Medical Sensor Networks.....</b>	<b>288</b>
<i>Md Zakirul Alam Bhuiyan, Mdaliuz Zaman, Guojun Wang, Tian Wang, and Jie Wu</i>	
<b>HealthEdge: Task Scheduling for Edge Computing with Health Emergency and Human Behavior Consideration in Smart Homes .....</b>	<b>290</b>
<i>Haoyu Wang, Jiaqi Gong, Yan Zhuang, Haiying Shen, and John Lach</i>	
<b>Author Index .....</b>	<b>293</b>
<b>Sponsors and Supporters .....</b>	<b>296</b>