

2024 IEEE 17th International Conference on Cloud Computing (CLOUD 2024)

**Shenzhen, China
7 – 13 July 2024**



**IEEE Catalog Number: CFP24CLO-POD
ISBN: 979-8-3503-6854-3**

**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:	CFP24CLO-POD
ISBN (Print-On-Demand):	979-8-3503-6854-3
ISBN (Online):	979-8-3503-6853-6
ISSN:	2159-6182

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

2024 IEEE 17th International Conference on Cloud Computing (CLOUD) **CLOUD 2024**

Table of Contents

Congress Steering Committee Chair Message	xii
Congress General Chair Message	xiii
Congress Program Chairs Message	xv
TCSVC Chair Message	xvi
CLOUD 2024 Chairs Message	xvii
CLOUD 2024 Committees	xviii
CLOUD 2024 Reviewers	xxii

CLOUD 1: Cloud and AI I

Fed2PKD: Bridging Model Diversity in Federated Learning via Two-Pronged Knowledge Distillation	1
<i>Zaipeng Xie (Hohai University), Han Xu (Hohai University), Xing Gao (Hohai University), Junchen Jiang (Hohai University), and Ruiqian Han (The Hong Kong University of Science and Technology (Guangzhou))</i>	
Resource Efficient Bayesian Optimization	12
<i>Namit Juneja (University at Buffalo, USA), Varun Chandola (University at Buffalo, USA), Jaroslaw Zola (University at Buffalo, USA), Olga Wodo (University at Buffalo, USA), and Parth Desai (University at Buffalo, USA)</i>	
Improving Federated Learning through Low-Entropy Client Sampling Based on Learned High-Level Features	20
<i>Waqwoya Abebe (Iowa State University, USA), Pablo Munoz (Intel Labs, USA), and Ali Jannesari (Iowa State University, USA)</i>	

CLOUD 2: Cloud-based Applications

Intent-Driven Multi-engine Observability Dataflows for Heterogeneous Geo-Distributed Clouds	30
<i>Pankaj Thorat (IBM Research India)</i>	
ArcaDB: A Disaggregated Query Engine for Heterogenous Computational Environments	42
<i>Kristalys Ruiz-Rohena (University of Puerto Rico, USA) and Manuel Rodriguez-Martinez (University of Puerto Rico, USA)</i>	

TraceMesh: Scalable and Streaming Sampling for Distributed Traces	54
<i>Zhuangbin Chen (Sun Yat-sen University, China), Zhihan Jiang (The Chinese University of Hong Kong, China), Yuxin Su (Sun Yat-sen University, China), Michael R. Lyu (The Chinese University of Hong Kong, China), and Zibin Zheng (Sun Yat-sen University, China)</i>	
Predictive Placement of Geo-distributed Blockchain Nodes for Performance Guarantee	66
<i>Junseok Lee (Korea University, Korea), Yeonho Yoo (Korea University, Korea), Chuck Yoo (Korea University, Korea), and Gyeongsik Yang (Korea University, Korea)</i>	

CLLOUD 3: Cloud Service Delivery I

Dynamic Workflow Scheduling in the Edge-Cloud Continuum: Optimizing Runtimes under Budget Constraints	69
<i>Stefan Pedratscher (University of Innsbruck, Austria), Thomas Fahringer (University of Innsbruck, Austria), and Juan Aznar-Poveda (University of Innsbruck, Austria)</i>	
FastMig: Leveraging FastFreeze to Establish Robust Service Liquidity in Cloud 2.0	81
<i>Sorawit Manatura (Kasetsart University, Thailand), Chantana Chantrapornchai (Kasetsart University, Thailand), Thanawat Chanikaphon (University of Louisiana at Lafayette, USA), and Salehi Mohsen Amini (University of North Texas, USA)</i>	
Telemetry-Driven Microservices Orchestration in Cloud-Edge Environments	91
<i>Angelo Marchese (University of Catania) and Orazio Tomarchio (University of Catania)</i>	
Rethinking Application Container Networking in a Multi-cluster world	102
<i>Chander Govindarajan (IBM Research, India), Priyanka Naik (IBM Research, India), Kavya Govindarajan (IBM Research, India), and Seep Goel (IBM Research, India)</i>	

CLLOUD 4: Cloud Infrastructure I

Optimizing Simultaneous Autoscaling for Serverless Cloud Computing	105
<i>Harold Ship (IBM Research - Israel), Evgeny Shindin (IBM Research - Israel), Chen Wang (IBM Research), Diana Arroyo (IBM Research), and Asser Tantawi (IBM Research)</i>	
UniNet: Accelerating the Container Network Data Plane in IaaS Clouds	115
<i>Yuan Ma (UIUC), Scott Smith (UIUC), Bill Dai (UIUC), Hubertus Franke (IBM Research), Bharat Sukhwani (IBM Research), Sameh Asaad (IBM Research), Jinjun Xiong (University at Buffalo), Volodymyr Kindratenko (UIUC), and Deming Chen (UIUC)</i>	
Availability-guarantee and Traffic Optimization Virtual Machine Placement in 5G Cloud Datacenters	128
<i>Wencong Yang (Zhengzhou University, School of Electrical Engineering), Shouyi Yang (Zhengzhou University, School of Electrical Engineering), Yi Yue (China Unicom Research Institute), Tian Chen (Zhengzhou University, School of Electrical Engineering), and Wanming Hao (Zhengzhou University, School of Electrical Engineering)</i>	

The Impact of Low-Entropy on Chunking Techniques for Data Deduplication	134
<i>Mu'men Al Jarah (University of Waterloo, Canada), Sreeharsha Udayashankar (University of Waterloo, Canada), Abdelrahman Baba (University of Waterloo, Canada), and Samer Al-Kiswany (University of Waterloo, Canada)</i>	

CLLOUD 5: Cloud Security I

CRISP: Confidentiality, Rollback, and Integrity Storage Protection for Confidential Cloud-Native Computing	141
<i>Ardhi Putra Pratama Hartono (TU Dresden), Andrey Brito (Universidade Federal de Campina Grande), and Christof Fetzer (TU Dresden)</i>	
Data Splitting based Double Layer Encryption for Secure Ciphertext Deduplication in Cloud Storage	153
<i>Xin Tang (University of International Relations, China) and Luchao Jin (University of International Relations, China)</i>	
Securing AI Inference in the Cloud: Is CPU-GPU Confidential Computing Ready?	164
<i>Apoorve Mohan (IBM T.J. Watson Research Center Yorktown Heights, New York, USA), Mengmei Ye (IBM T.J. Watson Research Center Yorktown Heights, New York, USA), Hubertus Franke (IBM T.J. Watson Research Center Yorktown Heights, New York, USA), Mudhakar Srivatsa (IBM T.J. Watson Research Center Yorktown Heights, New York, USA), Zhuoran Liu (IBM T.J. Watson Research Center Yorktown Heights, New York, USA), and Nelson Mimura Gonzalez (IBM T.J. Watson Research Center Yorktown Heights, New York, USA)</i>	
A Secure Framework for Continuous Compliance across Heterogeneous Policy Validation Points..	176
<i>Takumi Yanagawa (IBM Research - Tokyo), Vikas Agarwal (IBM Research - India), Yuji Watanabe (IBM Research - Tokyo), Lou Degenaro (IBM Thomas J. Watson Research Center), and Anca Sailer (IBM Thomas J. Watson Research Center)</i>	

CLLOUD 6: Cloud Management & Operations I

TempoScale: A Cloud Workloads Prediction Approach Integrating Short-Term and Long-Term Information	183
<i>Linfeng Wen (Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences, China, University of Chinese Academy of Sciences, China), Minxian Xu (Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences, China), Adel N. Toosi (Faculty of Information Technology, Monash University, Australia), and Kejiang Ye (Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences, China)</i>	
LGDCloudSim: A Resource Management Simulation System for Large-Scale Geographically Distributed Cloud Data Center Scenarios	194
<i>Jiawen Liu (Tongji University, China), Yuehao Xu (Tongji University, China), Binbin Feng (Tongji University, China), and Zhijun Ding (Tongji University, China; The Shanghai Artificial Intelligence Laboratory, China)</i>	

One System Call Hook to Rule All TEE OSes in the Cloud	205
<i>Kailun Qin (Shanghai Jiao Tong University, China) and Dawu Gu (Shanghai Jiao Tong University, China)</i>	
Optimizing Cloud Workloads: Autoscaling with Reinforcement Learning	217
<i>Pratik Mishra (IBM Research, India), Sandeep Hans (IBM Research, India), Diptikalyan Saha (IBM Research, India), and Pratibha Moogi (IBM Research, India)</i>	

CLOUD 7: Cloud Management & Operations II

Towards Sustainable Cloud Software Systems through Energy-Aware Code Smell Refactoring	223
<i>Asif Imran (California State University San Marcos), Tefik Kosar (University at Buffalo), Jaroslaw Zola (University at Buffalo), and M. Fatih Bulut (Microsoft)</i>	
Genetic Algorithm with Repair Method for Deadline-Constrained IoT Workflow Scheduling in Fog-Cloud Computing	235
<i>Amer Saeed (Victoria University of Wellington), Gang Chen (Victoria University of Wellington), Hui Ma (Victoria University of Wellington), and Qiang Fu (RMIT University)</i>	
Analysis and Optimization for Passive One-way Delay Measurement Tax in Container Networks .	247
<i>Jingzhao Xie (University of Electronic Science and Technology of China, China), Chongxi Ma (University of Electronic Science and Technology of China, China), Hongfang Yu (University of Electronic Science and Technology of China, China), Long Luo (University of Electronic Science and Technology of China, China), and Gang Sun (University of Electronic Science and Technology of China, China)</i>	
Syscall Analysis for Resource Stress Identification for Container Network Functions	256
<i>Chander Govindarajan (IBM Research, India), Priyanka Naik (IBM Research, India), Kavya Govindarajan (IBM Research, India), Seep Goel (IBM Research, India), Palanivel Kodeswaran (IBM Research, India), Sayandeep Sen (IBM Research, India), and Praveen Jayachandran (IBM Research, India)</i>	

CLOUD 8: Cloud and AI II

S2TAR: Shared Secure Trusted Accelerators with Reconfiguration for Machine Learning in the Cloud	267
<i>Wei Ren (University of Illinois Urbana-Champaign, USA), Sandhya Koteswara (IBM Research, USA), Mengmei Ye (IBM Research, USA), Hubertus Franke (IBM Research, USA), and Deming Chen (University of Illinois Urbana-Champaign, USA)</i>	
A Comparative Analysis of Generative Adversarial Networks for Generating Cloud Workloads	279
<i>Niloofer Sharifisadr (University of Calgary, Canada), Diwakar Krishnamurthy (University of Calgary, Canada), and Yasaman Amannejad (Mount Royal University, Canada)</i>	

SAM: Subseries Augmentation-based Meta-learning for Generalizing AIOps Models in Multi-cloud Migration	291
<i>Xi Yang (IBM Research), Paulito Palmes (IBM Research), Saurabh Jha (IBM Research), Bekir Turkkkan (IBM Research), Gerard Vanloo (IBM Research), Frank Bagehorn (IBM Software), Chandra Narayanaswami (IBM Research), Larisa Shwartz (IBM Consulting), Naoki Abe (IBM Research), Yu Deng (IBM Research), and Daby M. Sow (IBM Research)</i>	
Harmonia: Accurate Federated Learning with All-Inclusive Dataset	302
<i>Wonmi Choi (Korea University, Korea), Juyoung Ahn (Korea University, Korea), Yeonho Yoo (Korea University, Korea), Chuck Yoo (Korea University, Korea), and Gyeongsik Yang (Korea University, Korea)</i>	

CLOUD 9: Cloud-based Applications II

SimuScale: Optimizing Parameters for Autoscaling of Serverless Edge Functions through Co-Simulation	305
<i>Philipp Raith (TU Wien, Austria), Stefan Nastic (TU Wien, Austria), and Schahram Dustdar (TU Wien, Austria)</i>	
Distributed Dataflow Across the Edge-Cloud Continuum	316
<i>Tyler Ekairb (University of California, Santa Barbara), Lukas Brand (HAW Landshut), Nagarjun Avaraddy (University of California, Santa Barbara), Markus Mock (HAW Landshut), Chandra Krintz (University of California, Santa Barbara), and Rich Wolski (University of California, Santa Barbara)</i>	
A Technique for Secure Variant Calling on Human Genome Sequences Using SmartNICs	328
<i>Praveen Rao (University of Missouri, USA) and Khawar Shehzad (University of Missouri, USA)</i>	
Decoding Logs for Automatic Metric Identification	336
<i>Pranjal Gupta (IBM Research), Prateeti Mohapatra (IBM Research), Debanjana Kar (IBM Research), Seema Nagar (IBM Research), Jae-wook Ahn (IBM Research), Amit Paradkar (IBM Research), and Mudhakar Srivatsa (IBM Research)</i>	

CLOUD 12: Cloud Infrastructure II

SpotKV: Improving Read Throughput of KVS by I/O-aware Cache and Adaptive Cuckoo Filters ...	344
<i>Yi Liu (University of California Santa Cruz, USA), Ruilin Zhou (University of California Santa Cruz, USA), Yuhang Gan (University of California Santa Cruz, USA), and Chen Qian (University of California Santa Cruz, USA)</i>	
SlackAuto: Slack-Aware Vertical Autoscaling of CPU Resources for Serverless Computing	355
<i>Hyeon-Jun Jang (Konkuk University) and Hyun-Wook Jin (Konkuk University)</i>	
OS4C: An Open-Source SR-IOV System for SmartNIC-based Cloud Platforms	365
<i>Scott Smith (UIUC), Yuan Ma (UIUC), Marissa Lanz (UIUC), Bill Dai (UIUC), Martin Ohmacht (IBM Research), Bharat Sukhwani (IBM Research), Hubertus Franke (IBM Research), Volodymyr Kindratenko (UIUC), and Deming Chen (UIUC)</i>	

CLOUD 10: Cloud-based Applications III

Budget Aware Performance Testing of Microservices	376
<i>Quinn Cooper (University of Calgary, Canada), Diwakar Krishnamurthy (University of Calgary, Canada), and Yasaman Amannejad (University of Calgary, Canada)</i>	
Enabling Programmable Metric Flows	386
<i>Aishwariya Chakraborty (IBM Research), Chander Govindarajan (IBM Research), Kavya Govindarajan (IBM Research), Priyanka Naik (IBM Research), and Seep Goel (IBM Research)</i>	
Efficient Application Image Management in the Compute Continuum: A Vertex Cover Approach Based on the Think-Like-a-Vertex Paradigm	399
<i>Emanuele Carlini (ISTI-CNR), Patrizio Dazzi (University of Pisa), Antonios Makris (Harokopio University of Athens and National Technical University of Athens), Matteo Mordacchini (IIT-CNR), Theodoros Theodoropoulos (Harokopio University of Athens and National Technical University of Athens), and Konstantinos Tserpes (Harokopio University of Athens and National Technical University of Athens)</i>	
CodeSift: An LLM-Based Reference-Less Framework for Automatic Code Validation	404
<i>Pooja Aggarwal (IBM Research, India), Oishik Chatterjee (IBM Research, India), Ting Dai (IBM Research, USA), Prateeti Mohapatra (IBM Research, India), Brent Paulovicks (IBM Research, USA), Brad Blancett (IBM, USA), and Arthur De Magalhaes (IBM, USA)</i>	

CLOUD 13: Cloud Service Delivery II

Delay-Aware Service Caching in Edge Cloud: An Adversarial Semi-Bandits Learning-based Approach	411
<i>Jinpeng Li (Chongqing University, China), Yunni Xia (Chongqing University, China), Xiaoning Sun (Chongqing Normal University, China), Peng Chen (Xihua University, China), Xiaobo Li (Chongqing Animal Husbandry Techniques Extension Center, China), and Jiafeng Feng (China Huadian Corporation Guangdong Branch, China)</i>	
Coordinating Compaction between LSM-tree based Key-Value Stores for Edge Federation	419
<i>Jeeseob Kim (Sogang University, Republic of Korea), Honghyeon Yoo (Sogang University, Republic of Korea), Seungjae Lee (Sogang University, Republic of Korea), Hongsu Byun (Sogang University, Republic of Korea), and Sungyong Park (Sogang University, Republic of Korea)</i>	
The State of FaaS: An Analysis of Public Functions-as-a-Service providers	430
<i>Nnamdi Ekwe-Ekwe (University of St Andrews) and Lucas Amos (lucasamos.dev)</i>	
GASS: GPU Automated Sharing at Scale	439
<i>Jiafan Zhu (Google LLC, USA), Xiao Zhang (Google LLC, USA), Konstantinos Menychtas (Google LLC, USA), Zhijing Gene Qin (Google LLC, USA), Steven Hand (Google LLC, USA), Dragos Sbirlea (Google LLC, USA), and Yuang Liu (Google LLC, USA)</i>	

CLOUD 11: Cloud-based Applications IV

Optimizing Context Caching using a Novel Hybrid Strategy for Dynamically Monitoring Access Probability	446
<i>Ashish Manchanda (Swinburne University of Technology, Australia), Prem Prakash Jayaraman (Swinburne University of Technology, Australia), Abhik Banerjee (Swinburne University of Technology, Australia), and Arkady Zaslavsky (Deakin University, Australia)</i>	
Process-based Efficient Power Level Exporter	456
<i>Marcelo Amaral (IBM Research - Tokyo), Huamin Chen (Red Hat - USA), Tatsuhiro Chiba (IBM Research - Tokyo), Rina Nakazawa (IBM Research - Tokyo), Sunyanan Choochootkaew (IBM Research - Tokyo), Eun Kyung Lee (IBM Research - USA), and Tamar Eilam (IBM Research - USA)</i>	
iContinuum: An Emulation Toolkit for Intent-Based Computing Across the Edge-to-Cloud Continuum	468
<i>Negin Akbari (Monash University, Australia), Adel N. Toosi (Monash University, Australia), John Grundy (Monash University, Australia), Hourieh Khalajzadeh (Deakin University, Australia), Mohammad Sadegh Aslanpour (Monash University, Australia), and Shashikant Ilager (TU Wien, Austria)</i>	
DRLQ: A Deep Reinforcement Learning-based Task Placement for Quantum Cloud Computing ...	475
<i>Hoa T. Nguyen (The University of Melbourne, Australia), Muhammad Usman (The University of Melbourne and CSIRO, Australia), and Rajkumar Buyya (The University of Melbourne, Australia)</i>	

CLOUD 14: Cloud-based Applications V

Self Adjusting Log Observability for Cloud Native Applications	482
<i>Divya Pathak (Indian Institute of Technology Hyderabad, India), Mudit Verma (IBM Research, India), Aishwariya Chakraborty (IBM Research, India), and Harshit Kumar (IBM Research, India)</i>	
Carbon-Aware and Fault-Tolerant Migration of Deep Learning Workloads in the Geo-distributed Cloud	494
<i>Jeonghyeon Park (Dong-A University, Korea), Daero Kim (Dong-A University, Korea), Jiseon Kim (Dong-A University, Korea), Jungkyu Han (Dong-A University, Korea), and Sejin Chun (Dong-A University, Korea)</i>	
MediatorDNN: Contention Mitigation for Co-located DNN Inference Jobs	502
<i>Seyed Morteza Nabavinejad (Worcester Polytechnic Institute, USA), Sherief Reda (Brown University, USA), and Tian Guo (Worcester Polytechnic Institute, USA)</i>	
Author Index	513