2026 IEEE 46th International Conference on Distributed Computing Systems Workshops (ICDCSW 2026)

Seoul, Republic of Korea
22-25 June 2026

IEEE Catalog Number: CFP2628C-POD
ISBN: 979-8-3195-3648-8

IEEE
Advancing Technology for Humanity

Copyright © 2026, IEEE

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:	CFP2628C-POD
ISBN (Print-On-Demand):	979-8-3195-3648-8
ISBN (Online):	979-8-3195-3647-1
ISSN:	1545-0678

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

When CXL Meets Remote Atomic Operations, Cache Coherent NIC Demand Rearchitecturing 1
Zihao Wei, Dezun Dong, Lizhou Wu, Hanzhi Xun, Liquan Xiao, Zhen Ruan

A Multi-Head Attention Approach for SLA Compliance Monitoring in Data Centers 5
Omanshu Thapliyal

A Multi-Head Attention Approach for SLA Compliance Monitoring in Data Centers 11
Omanshu Thapliyal

DCDR: Device–Cloud Collaborative Framework for Transformer-Based Intelligent Signal Recognition 17
Guoliang Cheng, Wen Wu, Jinglin Huang, Jiayi Cong, Zhi Mao, Changsheng You

VSCI: Verifiable Smooth Coded Inference Via Spline Reconstruction 19
Ekaterina Pavlova, Alexey Frolov

Towards a Hierarchical Entanglement Access Control 21
Akihisa Takahashi, Yuhsaku Nakai, Yoshito Tobe

Deployment-Aware Causal Root Cause Analysis in Distributed Computing Continuum Systems 27
Zhenghao Wu, Ruihai Dong, Hadi Tabatabaee Malazi

Noise-Calibrated Lossy Compression for Privacy Leakage Reduction in Federated Learning 33
Shan Huang, Zhijing Ye, Zhaorui Zhang, Sheng Di, Jiamin Wang, Wendy Hui Wang, Xiaodong Yu

ADAS-LORA: Efficient Federated LLM Tuning in Edge Systems by Decoupling Compute and Network 39
Pablo Espinosa-Campos, Lucas Liebe, Thanh-Tung Nguyen, Nhat-Quang Tau, Yuheng Wu, Dongman Lee

SplitQ-PEFT: Quantum-Enhanced PEFT for Communication-Efficient Split Learning 45
Emily Jimin Roh, Soohyun Park, Joongheon Kim

Performance-Aware CPU/GPU Backend Selection and Code Adaptation for Python Scientific Computing in Heterogeneous Cloud Environments 49
Che Huang, Ayako Iwaki, Tomoya Hatano, Tatsuya Shimada

Distributed Multi-LLM Provenance Graph Construction for Scalable Correlation of APT Logs* 53
Jiwon Ock, Jungsik Lee, Changhee Choi

A Meta Stopping Mechanism for a KD-Based Decentralized System 57
Mbasa Joaquim Molo, Abednego Wamuhindo Kambale, Lucia Vadicamo, Claudio Gennaro Emanuele Carlini

GV-DRM: Gradient-Variance Driven Dynamic Rank Migration for Efficient LoRA Fine-Tuning in Federated Split Learning 61
Ali Irfan, Bong Jun Choi

Delay-Aware Robust Quantum Control for Overcoming Micro-Jitter in Modular Quantum Computing 65
Hoyeong Lee, Emily Jimin Roh, Joongheon Kim

Data-Free Early Stopping for Federated Learning with Consensus of Local Task Vectors 69
Youngjoon Lee, Hyukjoon Lee, Seungrok Jung, Andy Luo, Jinu Gong, Yang Cao, Joonhyuk Kang

Non-Redundant and Hotness-Aware Caching for Low-Latency Retrieval-Augmented Generation 73
Zeyu Xie, Hongtao Wang, Yu-Ang Zhang, Peiquan Jin

A Personalized Sparse LoRA Fine-Tuning Framework Under Federated Learning 77
Yinggang Sun, Wenting Zhang, Haining Yu, Xiangzhan Yu

Free-Rider Detection without Validation Data in Federated Learning Using Noise Injection 81
Geon Yong Yoon, Bong Jun Choi

Blackbird: A Scalable and Resource-Efficient Framework for Distributed Network Monitoring 85
Elisa Arnoldi, Gabriele Merli, Marco Abbadini, Michele Beretta, Dario Facchinetti, Matthew Rossi Stefano Paraboschi

Reliable Federated Fine-Tuning for Multilingual ASR 89
Shivani Sanjay Kolekar, Md Ilias Bappi, Kyungbaek Kim

ZeroTrustNet: A Software-Defined Perimeter Framework with Decentralised Identifiers for Zero Trust Network Security 93
Baptiste Rodrigues, Kilian Bonnau, Ounich Nizar, Lamine Amour, Sami Souihi

Fafnir: A Flexible Load Generator for Reproducible ML Inference Benchmarks 97
Theo Radig, Holger Karl, David Simon Nelle

An AI-Enhanced SSH Honeypot for Real-Time Threat Intelligence and Interactive Security Analytics 101
Samarth Devkar, Min Chen

Interactive Demonstration of Asynchronous Quantum-as-a-Service for Autonomous Driving 105
Emily Jimin Roh, Joongheon Kim

Ledgera and its Application to Decentralised Collaborative Fine-Tuning 109
Antonella Del Pozzo, Erwan Mahe, Sami Souihi, Sara Tucci-Piergiovanni

CAED: Context-Aware Data Cleaning Tool 113
Janus Roestenburg, Abdulhakim A. Qahtan

DySkew: Dynamic Data Redistribution for Skew- Resilient Snowpark UDF Execution 117
Chenwei Xie, Urjeet Shrestha, Corbin, Lukas Lorimer, Gopal V, Zihao Ye, Yi Pan, Nic Crouch Elliott Brossard, Florian Funke, Yuxiong He

MARL: A Multi-Agent Reinforcement Learning Framework for Buffer Management in Multi-Tenant Cloud Databases 123
Xin Huang, Yu-Ang Zhang, Hongtao Wang, Peiquan Jin

Predicate-Agnostic Filtered ANNS for Low-Selectivity Hybrid Queries in FAISS 129
Yichen Li, Junli Wang, Peiquan Jin, Xiaoniu Fang

HAP^5: Heterogeneity-Aware Pipeline Parallelism Planning and Performance Prediction 135
Anurag Dutt, Buvana Ramanan, Brian Friedman, Anshul Gandhi, Manzoor Khan

Large Language Models as a Stress Test for Enterprise API Gateways 141
Zixin Wang

Silent Data Corruption in Distributed Genai-Augmented Data Pipelines: An Empirical Study of Hallucination-Induced Failure Propagation 147
Chalapathi Koneni, Sanjay Lokula

Llm-Cache: an Efficient Context-Aware Semantic Caching Framework for Distributed Llm Inference Services 153
Haoying Jin, Haoyang Feng

Scene Complexity-Aware Adaptive Batch Separation for Robust Camera-based BEV Perception 159
Yong Chan Lee, Eun-Kyu Lee, Jiwoo Lee

Silent Data Corruption in Distributed Genai-Augmented Data Pipelines: an Empirical Study of Hallucination-Induced Failure Propagation 165
Chalapathi Koneni, Sanjay Lokula

Decentralized and AI-Enhanced IoE Framework: P2P Bartering, Donation, and On-Chain Learning 172
Abdulrezzak Zekiye, Öznur Özkasap, Ouns Bouachir, Moayad Aloqaily

AI Agent-Driven Closed-Loop Control for Integrated Sensing and Communication Systems 174
Xinghan Cao, Wen Wu, Liang Li, Changyang She

Fleet-Level Optimization for Heterogeneous Device Provisioning in ML Model Deployments 176
Harry H. Jiang, Carlee Joe-Wong

A Hardware-Software Co-Design Framework for Spatial and Temporal Optimization in Distributed Quantum Architectures 178
Raymond P. H. Wu

Distributed Quantum Circuit Optimisation: Evaluating Global and Local Encodings 180
Maria Gragera Garces, Majid Haghparast

On the Distortion of Partitioning Performance by Random Quantum Circuits 184
Maria Gragera Garces

HPC-VQPU: Exporting Device-Aware Quantum Simulation From Batch-Scheduled HPC as an Interactive Service 190
Shusen Liu, Pascal Jahan Elahi, Ugo Varatto

Distributed Variational Quantum Optimisation by Entanglement-Selective Transport 196
Edric Matwiejew, Pascal Elahi, Ugo Varetto

Emulating NetQMPI Applications with CUNQA: A Decoupled Architecture for HPC Environments 202
Jorge Vázquez-Pérez, F. Javier Cardama, Tomás F. Pena, Andrés Gómez

Multi-Agent Proximal Policy Optimization with Dual Queues for Dynamic Task Offloading in Mobile Edge Computing 208
Zixuan Wang, Mingzhou Yang

Edge-Native Environmental Sensor Interpretation Using Lightweight Large Language Models 214
Dires Negash Fente, Cássio V. S. Prazeres

SpectralShard: Lightweight Partitioning Under Dynamic Interaction Patterns 220
Yujin Pang, Clément Burgelin, Antoine Murat, Rachid Guerraoui

Continuum-Aware End-Cloud Offloading for Collaborative Inference 226
Shangyang Wu, Xingwei Wang, Yuhai Zhao, Yingyou Wen, Zhijiang Wang, Cong Wei

Intent-Based Web Application for Managing Sdn Research Clusters 232
Walid Nouicer, Jani Launonen, Lauri Lovén, Dinesh Sah

Quantifying Inter-City Network Latency in Europe: A Measurement based Study for Time-Critical Cloud Services 238
Thomas Schleicher, Kurt Horvath, Dragi Kimovski, Bernd Spiess, Oliver Hohlfeld

Fair Comparison of Scheduling Algorithms on Heterogeneous Edge Clusters: A Continuous Adaptive Benchmark 244
Zihang Wang, Boris Sedlak, Juan Luis Herrera, Schahram Dustdar

Federated Digital Twin-Driven Foundation Models for Energy-Aware Scheduling in Edge Data Centers 250
Andy Yuen-Khai Khoo, Yen-Hsun Meng, Cheng-Hou Chou, Chia-Hsuan Wu, Mu-Chun Su, Yi-Zeng Hsieh

A Communication Cost Model for Federated Mask-Voting 256
Saul Urso, Emanuele Carlini, Patrizio Dazzi, Matteo Mordacchini