ACM International Conference on Resilience, AI and Cyber-Systems (ACM RACS 2026)

November 17–20, 2026
Fukuoka, Japan

Final Program



SIGAPP

Organizing Committee

Conference Chair
Shih-Hao Hung, National Taiwan University, Taiwan

Program Chairs
Ya-Shu Chen, National Taiwan University of Science and Technology, Taiwan
Junggab Son, University of Nevada Las Vegas, USA

Poster Chair
Austin Hanson, Versaterm Corporation, USA

Publication Chair
Kwanghee Won, South Dakota State University, USA

Financial Chair
Hong Min, Gachon University, Korea

Registration & Website Chair
Joonhyouk Jang, Hannam University, Korea

Workshop Chair
Chulwoo Pack, South Dakota State University, USA

Publicity Chairs
Zhongwen Guo, Ocean University of China, China
Hyung-Do Choi, ETRI, Korea
Christian Esposito, University of Salerno, Italy
Chang Choi, Gachon University, Korea
Junyoung Heo, Hansung University, Korea
Chi-Sheng Shih, National Taiwan University, Taiwan
Wei Wang, San Diego State University, CA, USA
Jun Huang, South Dakota State University, USA
Chia-Heng Tu, National Cheng Kung University, Taiwan

Steering Committee

Sung Y. Shin, South Dakota State University, USA
Chih-Cheng Hung, Kennesaw State University, USA
Tei-Wei Kuo, Delta Electronics and National Taiwan University, Taiwan
John Kim, Utica University, USA
Juw Won Park, University of Louisville, KY, USA
Jiman Hong, Soongsil University, Korea

* Cover image generated using AI

Conference Schedule

Tuesday (November 17)		
14:00 – 16:30	Organizing Committee Meeting	TBD
16:30 – 18:00	RACS Business Meeting	TBD
Wednesday (November 18)		
Registration is available from Wednesday to Thursday		
09:30 – 11:00	Opening and Keynote	10A
11:00 – 11:30	Coffee Break	10A
11:30 – 13:00	Technical Session 1	10B
	Technical Session 2	10C
13:00 – 14:30	Luncheon	10A
14:30 – 16:00	Poster Session	10A
16:00 – 16:30	Coffee Break	10A
16:30 – 17:30	SIGAPP Meeting	TBD
18:30 – 20:30	Welcome Reception	TKP Garden City
Thursday (November 19)		
09:30 – 11:00	Technical Session 3	10B
	Technical Session 4	10C
11:00 – 11:30	Coffee Break	10A
11:30 – 13:00	Technical Session 5	10B
	Technical Session 6	10C
13:00 – 14:30	Luncheon	10A
14:30 – 17:00	Technical Session 7	10B
	Technical Session 8	10C
19:00 – 21:00	ACM RACS Banquet	The Kego Club by Happo-En (9th FL)
Friday (November 20)		
10:00 – 12:00	Future RACS Organization Meeting	TBD
13:30 – 14:30	RACS Wrap-up Meeting	TBD
15:00 – 16:30	Steering Committee Meeting	TBD
17:30 – 19:00	RACS Business Meeting	TBD

Preface

With the expansion of advanced computing and network technologies, innovative solutions are developed to address complex problems in all areas of industry and science. These efforts increasingly focus on resilient computing, artificial intelligence, and cyber-systems, encompassing fields such as programming, software engineering, graphics, databases, wireless networks, security, distributed systems, operating systems, quantum computing, and more. The International Conference on Resilience, AI and Cyber-Systems (RACS) provides a forum for exchanging highly original ideas about systematic, dynamic, practical solutions and aims primarily at researchers who have experience in designing such computing systems and applications. Each year, RACS brings together engineers and scientists from diverse communities with interests in practical computing technologies and creates an environment for them to exchange experimental results, novel designs, works-in-progress, experiences, case studies, and trend-setting ideas.

This volume contains the papers selected for presentation at RACS 2026, which is held in Fukuoka, Japan, November 17–20, 2026. There were 86 submissions from 11 countries. Each paper underwent a blind peer review process by members of the Technical Program Committee (TPC), and 32 regular full papers and 13 poster papers were accepted for publication in the proceedings of the conference. The selected papers cover a wide range of topics, including innovative theories, technologies, and applications that contribute to the design and operation of resilient, intelligent, and secure cyber-systems.

The RACS 2026 Technical Program Committee was led by two Program Chairs: Dr. Ya-Shu Chen and Dr. Junggab Son. We wish to thank the TPC members for their valuable time and technical input in creating an excellent program. Without their help and contributions, this conference would not have been possible. We would also like to thank the Organizing Committee for their great effort. Our special thanks go to Dr. Kwanghee Won (Publication Chair), Dr. Hong Min (Financial Chair), Dr. Austin Hanson (Poster Chair), Dr. Joonhyouk Jang (Registration & Website Chair), Dr. Zhongwen Guo, Dr. Hyung-Do Choi, Dr. Christian Esposito, Dr. Chang Choi, Dr. Junyoung Heo, Dr. Chi-Sheng Shih, Dr. Wei Wang, Dr. Jun Huang, Dr. Chia-Heng Tu (Publicity Chairs), and Dr. Chulwoo Pack (Workshop Chair) for their leadership and excellent work to organize the conference. Finally, we thank all of the authors who contributed to this volume and all the attendees of RACS 2026.

Dr. Shih-Hao Hung
Conference Chair, RACS 2026

Keynote Speech

Title: Unlocking the Full Potential of 3D NAND Flash



Speaker:
Chun Jason Xue, Ph.D.
Division of Computing and Mathematical Sciences
Mohamed Bin Zayed University of Artificial Intelligence
Abu Dhabi, United Arab Emirates

Abstract

This talk will present a decade of research on unlocking the full potential of 3D NAND flash storage through cross-layer optimization. As flash storage now dominates data center capacity, the industry faces a fundamental scaling trilemma where higher density simultaneously degrades reliability, performance, and lifetime. The research path begins with deep physical characterization, establishing that accurate understanding of device behavior, including retention, read disturb, and wear, must be the foundation of any optimization. Building on these insights, the work progresses to read performance optimization, exploiting error asymmetry and using designated cells as reliability indicators to achieve near-zero read retry rates. The research then demonstrates that smart data encoding can transform how data interacts with flash physics, turning invalidated data and entropy-aware coding into tools for extending device lifetime. We also challenge the long-held constraint of sequential page programming, showing that strategic reprogramming can fundamentally reshape the write path for significant latency gains. Throughout this progression, the unifying principle is cross-layer co-design: by allowing information to flow between device physics, flash chips, controllers, and applications, it becomes possible to push the Pareto frontier on reliability, performance, and lifetime simultaneously rather than trading one for another. Most recently, we explore the usage of Flash memory in the memory hierarchy of ML inference systems. A preliminary study on HBF will be presented.

About the Speaker

Prof. Chun Jason Xue is currently a professor of computer science at MBZUAI, Abu Dhabi. His research focuses on memory and storage systems. He is currently an associate editor for ACM Transactions on Embedded Computing Systems, ACM Transactions on CPS, and ACM Transactions on Storage. He is a distinguished member of ACM, and a fellow of IEEE.

Research Interests

Professor Xue's research focuses on systems and memories and the consideration of efficiency, reliability, and scalability. His recent work explores cross-layer optimization across storage devices, computer architecture, operating systems, runtimes, and applications. This includes improving the reliability, lifetime, and performance of flash memory and SSDs; developing high-

performance I/O and memory-management mechanisms; and rethinking data movement across storage, host memory, and GPUs. He is interested in how OS kernel evolution can be benefited from the latest LLMs.

Homepage: https://mbzuai.ac.ae/study/faculty/chun-jason-xue/

Wednesday (November 18)

(S.1) Storage and I/O Systems 11:30 – 13:00

Session Chair: Austin Hanson

1. *Yi-Shen Chen and Chia-Hao Chang,* ***PAZA: Parallel-Aware Zone Allocation for Improving LSM-Tree Performance on ZNS SSDs***
2. *Yi-Han Lien, Dong-Xian Wu, Yen-Ting Chen, Yu-Pei Liang, Wei-Kuan Shih, and Yuan-Hao Chang,* ***Tier-aware Track Allocation Scheme for LSM-tree on Interlaced Magnetic Recording Hard Disk Drives***
3. *Chun-Ting Lai, Jen-Wei Hsieh, and Hung-Pin Chen,* ***FreshGC: Avoid Stale Free Blocks in 3D NAND Flash Memory***
4. *Guan-Tze Li, Liang-Chi Chen, Chien-Chung Ho, Yu-Ming Chang, and Yuan-Hao Chang,* ***Bi-Tiered Page Caching Design over CXL-enabled SSDs for Unlocking Fine-Grained I/O Bottleneck***

(S.2) Trustworthy Federated and Agentic AI 11:30 – 13:00

Session Chair: Ya-Shu Chen

1. *Marco De Santis, Aiman Faiz, Christiancarmine Esposito, and Giovanna Truda,* ***RF-ADAS: Responsible Federated Learning for Auditable Driver Monitoring in Advanced Driver-Assistance Systems***
2. *Tzu-En Yang, Yi-Chen Huang, Chih-Wei Chen, Tsen-Chih Hsu, Shu-Ping Lu, and Yung-Feng Lu,* ***Outcome-Guided Trust Governance for Secret-Shared Federated Learning under Secure-Aggregation Visibility Constraints***
3. *Pei-Yun Chang and Shih-Hao Hung,* ***Self-Preference Bias in LLM-as-a-Judge for Hardware Verification: Mechanisms and Training-Free Mitigation***
4. *Alexandros Konios, Nandhakumar Manikandasamu, and Bharathi S****, Adversarial Behavioral Mimicry: Length Normalization Breaks Behavioral Attestation for AI Agents***

Poster Session 14:30 – 16:00

1. *Hayoung Boo, Da-bin Im, Seong-Min Kim, and Junghoon Lee,* ***CNN-based Accident Trip Classification on a Mobility Platform***
2. *Da-bin Im, Seong-Min Kim, Hayoung Boo, and Junghoon Lee,* ***Design of a Rescoring Mechanism for Reducing False Alarms in Log-Based Crash Analysis***
3. *Chin-Fu Kuo, Guang-Chen Lin, Yung-Feng Lu, and Shu-Ping Lu,* ***SMGFL: Synchronous Multi-Group Federated Learning with Label-Flipping Attack Mitigation***
4. *Juchan Park, Joonhyouk Jang, Daehyeon Shin, Taehan Kim, Seongmin Park, Geumsook Heo, and Jiho Kang,* ***Interpreter-Based SPARC Emulation Framework***
5. *Franco Cirillo, Danial Daniyaruly, Christian Esposito, and Alina Gutoreva,* ***Modeling Cognitive-Load-Aware Phishing Defense through Agent-Based Modeling and Evolutionary Game Theory***
6. *Yu-Xin Chen, Tzung-Dau Wang, and Shih-Hao Hung,* ***A Cross-modal CAG–CCTA Stenosis Mapping Framework with Anatomical Feature Validation***
7. *Chi-Sheng Daniel Shih, Jong-Ruey Yang, Chi-Wei Chen, and Xin-You Liu,* ***CIMDiff: A CIM-Friendly UNet for Diffusion Models***
8. *Yewon Cha and Hong Min,* ***Attention Weights as Discriminative Signals for Wearable Walking-Risk Sub-classification***
9. *Abdul Hai Karimi, Saebom Lee, and Chang Choi,* ***DeepStegoChain: Blockchain-Assisted Adversarial Neural Network for Enhanced Image Steganography***
10. *Ziyang Gong, Namgyu Jung, and Chang Choi,* ***Emotion Recognition Using Multi-modal Physiological Signals***
11. *Jun-Yeong Kang, Yousung Yeon, Jeong-Ho Kim, and Chang Choi,* ***Degradation-Conditioned Alignment for Compression-Robust Deepfake Detection***
12. *Taimoor Khan, Junho Yoon, and Chang Choi,* ***BA-FedNet: Boundary-Aware and Domain-Adaptive Federated Learning for Heterogeneous Land Cover Classification***
13. *Chan Sik Han and Keon Myung Lee,* ***Evaluating Temporal Order Sensitivity in Regression Test-Time Adaptation***

Thursday (November 19)

(S.3) Emerging Memory and Data-Centric Architecture 09:30 – 11:00

Session Chair: Junyoung Heo

1. *Han-Yu Liao, Liang-Wei Lai, and Jen-Wei Hsieh,* ***Characterizing Frequency-Aware Placement for Skyrmion Racetrack Memory in Block Storage***
2. *Yu-Pei Liang, Jian-Yi Pan, Yen-Ting Chen, Shuo-Han Chen, Xin Liu, Wei-Kuan Shih, and Yuan-Hao Chang,* ***A Dual-Mode Skyrmion Racetrack Memory Architecture Bridging Word-Based and Bit-Interleaved Access for In-Memory Hyperdimensional Computing***
3. *Xin Liu, Yun-Shan Hsieh, Po-Chun Huang, Jonathan Huang, Wei-Kuan Shih, and Yuan-Hao Chang,* ***VASE-HD: Variance-Aware Summary Encoding for Hyperdimensional Computing***
4. *Jinheng Li, Guosheng Wang, Qiao Li, and Chun Jason Xue,* ***HBFSim: A Simulator for Exploring HBF in Memory Hierarchies for LLMs***

(S.4) Cybersecurity, Zero Trust, and Resilient Monitoring 09:30 – 11:00

Session Chair: Chia-Yin Liu

1. *Chun-Hao Huang, Chin-Fu Kuo, Shu-Ping Lu, and Yung-Feng Lu,* ***Post-Quantum Zero Trust Architecture for EHR Access Control: A DATES-Driven Approach***
2. *Alexander Gimourginas and Shajina Anand,* ***A Three-Layer Composite Risk Model for Endpoint Security Monitoring in K-12 Schools***
3. *Ren-Jie Chen, Shu-Ye Wang, and Shanq-Jang Ruan,* ***EdgeGuard: Stateful Event-Driven Monitoring and Agent-Assisted Response for Edge Physical Security***

(S.5) Multimodal Retrieval and Intelligent Perception 11:30 – 13:00

Session Chair: Shuo-Han Chen

1. *Abdolhossein Erfanshekooh, Kwanghee Won, John Kim, and Chulwoo Pack,* ***Content-Aware Page Promotion for Multimodal Document Retrieval-Augmented Question Answering***
2. *Sugam Mishra, Kwanghee Won, and Chulwoo Pack,* ***Semantic Think-on-Graph (SemToG): Enhancing GraphRAG through Semantic Community Detection***
3. *Chi-Sheng Daniel Shih, Ting-Fang Tsai, and Xiaodong Liu,* ***Q-CART: Integer-Only Progressive Quantization of a Color-Aided Point Cloud Registration Transformer***

(S.6) Heterogeneous Platforms and Real-Time Scheduling 11:30 – 13:00

Session Chair: Ing-Jer Huang

1. *Yu-Hui Wu, Xie-Ru Li, Po-Hsuan Huang, Teik-Hui Lee, Chia-Heng Tu, and Shih-Hao Hung,* ***PIQ: A Pluggable Interface for Quantum-Classical Execution Backends for Variational Quantum Algorithms***
2. *Jun-Soo Lee, Hong-Geun Kim, Junyoung Heo, and Hong Min,* ***Machine Learning Model for Predicting Execution Time of Deep Neural Networks on Diverse Hardware Platform***
3. *Shiwu Lo, Hsiao-Yueh Tsai, and Ting-Chia Kuo,* ***Design of Mobile Camera Systems After 1-Inch Image Sensors***
4. *Zong-Lun Zhang, Yu-Ting Huang, Yu-Rou Shen, and Yu-Pei Liang,* ***Priority-Aware Cluster Scheduling for Multicore Mixed-Criticality Systems***

(S.7) Efficient AI and LLM Systems 14:30 – 17:00

Session Chair: Yu-Pei Liang

1. *En-Ming Huang, Yao-Ting Hsieh, Hsiang-Yu Tsou, Mu-Chi Chen, Shih-Hao Hung, and Hsiang-Tsung Kung,* ***Towards Training Private LLMs: Exploring Fine-Tuning Language Models on Apple Silicon with RDMA over Thunderbolt***
2. *Yun Chi Chen, Yi-Chuan Liang, and Shih-Hao Hung,* ***Privacy–Latency Trade-offs for Edge–Cloud Split Inference of Large Language Models***
3. *Sri Sai Rakesh Nakkilla, Xiangyu Zhou, Chia-Yin Liu, Raghu Garikapaty, Pi-Cheng Hsiu, Jingtong Hu, and Jeff Zhang,* ***ColdLLM: A Runtime Thermal Controller for Sustainable LLM Inference at the Edge***
4. *William Dai, Dheeraj Joshi, Karthik Chandranna, Mark Canning, Jeff Zhang, and Yong Liu,* ***Bottleneck Killer: Control-Plane Adaptive Serving Plans for Production Recommendation Inference***
5. *Chi-Sheng Daniel Shih and Yuan-Chun Lo,* ***GO-CART: GPU Orchestration for Concurrent Heterogeneous LLM Serving via Deep Reinforcement Learning***

(S.8) Edge, IoT, and Networked Systems 14:30 – 17:00

Session Chair: Juw Won Park

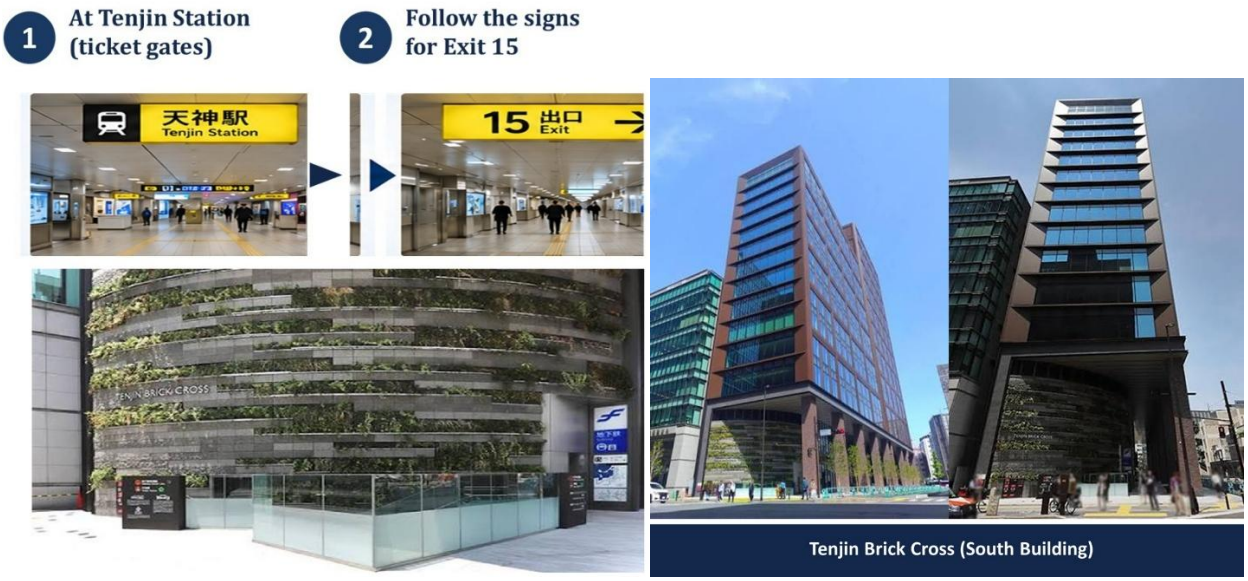
1. *Jihwan Yeom, Jinman Kim, and Joongjin Kook,* ***Energy-Aware PM2.5-Safe Multi-Agent AI Control for Synchronized Indoor–Outdoor Air-Quality Management***
2. *Yung-Feng Lu, Hung-Ming Chen, Shu-Ping Lu, and Ti-Wei Huang,* ***An Automated and Self-Healing Microservice Deployment System Across Cross-Domain Edge Environments Using KubeEdge***
3. *Hsueh-Wen Tseng, Ruei-Yu Wu, and Cheng-Huan Tai,* ***A Microburst-Detection-Based In-band Network Telemetry Framework on P4 Switches***
4. *Joohwan Park, Su Kim, Sunghyun Kim, and Dongkyun Kim,* ***SRG-cMAB: Contextual Multi-Armed Bandits for Coordinated Spatial Reuse***
5. *Yung-Feng Lu and Chen Han Qi,* ***Efficient RFID Duplicate Elimination: A Cloud-Edge Collaborative Signature Grid***

Venue

TKP Garden City Premium Tenjin (10th FL)

Japan, 〒810-0001 Fukuoka, Chuo Ward, Tenjin, 1 Chome−14−18 ブリッククロス南棟 10 階

Google map: [TKP Garden City PREMIUM Tenjin - Google Maps]



After going through the ticket gates of Tenjin Station (Fukuoka City Subway Airport Line), follow the "Exit 15 (出口 15)" signs. Tenjin Brick Cross (south building) is directly accessible through the underground passage.

Banquet

The Kego Club by Happo-En (9th FL)

Japan, 〒810-0001 Fukuoka, Chuo Ward, Tenjin, 2 Chome−2−20 警固神社社務所ビル 9 階

Google map: THE KEGO CLUB by HAPPO-EN - Google Maps



Approximately 5-minute walk from Tenjin Station (Fukuoka City Subway Airport Line). From Tenjin Station, follow the signs to Exit 4 (出口 4). After coming up to the ground level, walk toward Watanabe-dori Avenue. The building (9th FL) is a short walk from the exit.