

# **2025 IEEE International Conference on Cluster Computing Workshops (CLUSTER Workshops 2025)**

**Edinburgh, United Kingdom  
2-5 September 2025**



**IEEE Catalog Number: CFP2587K-POD  
ISBN: 979-8-3315-1257-6**

**Copyright © 2025 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:	CFP2587K-POD
ISBN (Print-On-Demand):	979-8-3315-1257-6
ISBN (Online):	979-8-3315-1256-9

**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

# Contents

<b>Welcome Message from the General Chairs</b>	<b>3</b>
<b>Welcome Message from the Workshop Chair</b>	<b>4</b>
<b>Welcome Message from the REX-IO Workshop Chairs</b>	<b>5</b>
<b>Welcome Message from LLMxHPC Workshop Chairs</b>	<b>7</b>
<b>Welcome Message from the Poster Committee Chair</b>	<b>9</b>
<b>Papers</b>	<b>11</b>
SYCL for CPU+GPU Heterogeneous Computing: A Study on Integrated and Discrete GPUs . . . . .	11
PALLAS: HPC trace analysis at scale . . . . .	13
Optimizing the Execution Core for Value Prediction in Superscalar Processors . . . . .	15
Parallel Counting of Non-Induced Graphlet Orbits . . . . .	17
SYCL QPU: an LLVM-based QPU simulation framework built using DPC++ . . . . .	19
NetQMPI: a practical MPI-inspired library for distributed quantum computing over NetQASM SDK . . . . .	21
Hyperspectral Imaging at Scale with Parallel DNN-based Regularization in Radio Astronomy . . . . .	23
FastEM: an efficient EM algorithm for learning Gaussian mixture models on compute clusters . . . . .	25
Hardware Emulation Framework for HPC Network . . . . .	27
Performance Characterization of Trinetra-A HPC Network Using MPI Benchmarks . . . . .	29
Performance Evaluation of HPC Benchmarks on a RISC-V Based Cluster	31
A similarity-aware MOE-based method for optimizing tensor programs across diverse GPUs . . . . .	33
Incremental Sparse Tensor Format for Maximizing Efficiency in Tensor-Vector Multiplications . . . . .	35
A Policy-Driven Approach for Securing Microservices Workflow in Kubernetes Cluster . . . . .	37
GPU-CPU Shared Memory Performance Analysis on NVIDIA GH200	39
Toward LLM-Compatible Log Representation Learning: A Hierarchical Semantic-Structural Framework for HPC Anomaly Detection	41
BBView: A View-Aware Burst-Buffer Mechanism for MPI-IO . . . . .	46
Benchmarking Darshan and Recorder for HPC I/O Profiling and Tracing	55
Optimizing I/O for an Exascale Implicit Kinetic Plasma Simulation using the Rabbit Storage System . . . . .	61
Towards an Optimal IO500 Configuration: Literature Meets Empirical Evaluation . . . . .	67