

2024 IEEE Symposium on High-Performance Interconnects (HOTI 2024)

**Albuquerque, New Mexico, USA
21 – 23 August 2024**



**IEEE Catalog Number: CFP24HIS-POD
ISBN: 979-8-3503-5602-1**

**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:	CFP24HIS-POD
ISBN (Print-On-Demand):	979-8-3503-5602-1
ISBN (Online):	979-8-3503-5601-4
ISSN:	1550-4794

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 Symposium on High-Performance Interconnects (HOTI) **HOTI 2024**

Table of Contents

Message from the General Chairs	vii
Message from the TPC Chairs	ix
Conference Committee	x
Keynote Speakers	xiii
Technical Sessions	xvi
Panel	xvii
Tutorials	xviii
Patrons and Sponsors	xxviii

Technical Paper Session A: Networks for Large Language Models

Rail-Only: A Low-Cost High-Performance Network for Training LLMs with Trillion Parameters	1
<i>Weiyang Wang (MIT CSAIL), Manya Ghobadi (MIT CSAIL), Kayvon Shakeri (Meta), Ying Zhang (Meta), and Naader Hasani (Meta)</i>	
Characterizing Communication in Distributed Parameter-Efficient Fine-Tuning for Large Language Models	11
<i>Nawras Alnaasan (The Ohio State University, USA), Horng-Ruey Huang (The Ohio State University, USA), Aamir Shafi (The Ohio State University, USA), Hari Subramoni (The Ohio State University, USA), and Dhabaleswar K. Panda (The Ohio State University, USA)</i>	

Technical Paper Session B: Interconnection Networks

Quality-of-Service Provision for BXI3-Based Interconnection Networks	20
<i>Miguel Sánchez de la Rosa (Department of Computing Systems, Universidad de Castilla-La Mancha, Spain), Gabriel Gomez-Lopez (Department of Computing Systems, Universidad de Castilla-La Mancha, Spain), Francisco J. Andújar (Departamento de Informática, Universidad de Valladolid, Spain), Jesus Escudero-Sahuquillo (Department of Computing Systems, Universidad de Castilla-La Mancha, Spain), José L. Sánchez (Department of Computing Systems, Universidad de Castilla-La Mancha, Spain), Francisco J. Alfaro (Department of Computing Systems, Universidad de Castilla-La Mancha, Spain), and Pierre-Axel Lagàdec (Atos, France)</i>	

Best Academic Paper Award: A New Mechanism to Identify Congesting Packets in High-Performance Interconnection Networks	24
<i>Cristina Olmedilla (Universidad de Castilla-La Mancha, Spain), Jesus Escudero-Sahuquillo (Universidad de Castilla-La Mancha, Spain), Pedro Javier Garcia (Universidad de Castilla-La Mancha, Spain), Francisco J. Quiles (Universidad de Castilla-La Mancha, Spain), Wenhao Sun (Huawei Technologies Co., Ltd., China), Long Yan (Huawei Technologies Co., Ltd., China), Yunping Lyu (Huawei Technologies Co., Ltd., China), and Jose Duato (Royal Spanish Academy of Sciences, Spain)</i>	

Technical Paper Session C: Collectives for Deep Learning

Towards a Standardized Representation for Deep Learning Collective Algorithms	33
<i>Jinsun Yoo (Georgia Institute of Technology), William Won (Georgia Institute of Technology), Meghan Cowan (NVIDIA), Nan Jiang (NVIDIA), Benjamin Klenk (NVIDIA), Srinivas Sridharan (NVIDIA), and Tushar Krishna (Georgia Institute of Technology)</i>	
Best Industry Paper Award: Unified Collective Communication (UCC): An Unified Library for CPU, GPU, and DPU Collectives	37
<i>Manjunath Gorentla Venkata (NVIDIA Corporation), Valentine Petrov (NVIDIA Corporation), Sergey Lebedev (NVIDIA Corporation), Devendar Bureddy (NVIDIA Corporation), Ferrol Aderholdt (NVIDIA Corporation), Joshua Ladd (NVIDIA Corporation), Gil Bloch (NVIDIA Corporation), Mike Dubman (NVIDIA Corporation), and Gilad Shainer (NVIDIA Corporation)</i>	

Technical Paper Session D: Optimizing Collective Operations

OHIO: Improving RDMA Network Scalability in MPI_Alltoall through Optimized Hierarchical and Intra/Inter-Node Communication Overlap Design	47
<i>Tu Tran (The Ohio State University, USA), Goutham Kalikrishna Reddy Kuncham (The Ohio State University, USA), Bharath Ramesh (The Ohio State University, USA), Shulei Xu (The Ohio State University, USA), Hari Subramoni (The Ohio State University, USA), Mustafa Abduljabbar (The Ohio State University, USA), and Dhabaleswar K. Panda (The Ohio State University, USA)</i>	
Demystifying the Communication Characteristics for Distributed Transformer Models	57
<i>Quentin Anthony (The Ohio State University, USA), Benjamin Michalowicz (The Ohio State University, USA), Jacob Hatef (The Ohio State University, USA), Lang Xu (The Ohio State University, USA), Mustafa Abduljabbar (The Ohio State University, USA), Aamir Shafi (The Ohio State University, USA), Hari Subramoni (The Ohio State University, USA), and Dhabaleswar K. Panda (The Ohio State University, USA)</i>	

Author Index	67
---------------------------	-----------