

2019 IEEE Symposium on High-Performance Interconnects (HOTI 2019)

**Santa Clara, California, USA
14 – 16 August 2019**



IEEE Catalog Number: CFP19HIS-POD
ISBN: 978-1-7281-5526-5

**Copyright © 2019 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:	CFP19HIS-POD
ISBN (Print-On-Demand):	978-1-7281-5526-5
ISBN (Online):	978-1-7281-5525-8
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

2019 IEEE Symposium on High- Performance Interconnects (HOTI) **HOTI 2019**

Table of Contents

Message from the General Chairs	viii
Message from the Technical Program Chairs	ix
Committees	x
Invited Talks	xii
Panel	xv
Tutorials	xvi
Keynotes	xxi
Sponsors	xxiii

Session A: Topologies

The First Supercomputer with HyperX Topology: A Viable Alternative to Fat-Trees?	1
<i>Jens Domke (RIKEN Center for Computational Science), Satoshi Matsuoka (RIKEN Center for Computational Science), Ivan Radanov Ivanov (Tokyo Institute of Technology), Yuki Tsushima (Tokyo Institute of Technology), Tomoya Yuki (Tokyo Institute of Technology), Akihiro Nomura (Tokyo Institute of Technology), Shin'ichi Miura (Tokyo Institute of Technology), Nic McDonald (Hewlett Packard Enterprise), Dennis Lee Floyd (Hewlett Packard Enterprise), and Nicolas Dubé (Hewlett Packard Enterprise)</i>	
Path2SL: Optimizing Head-of-Line Blocking Reduction in InfiniBand-Based Fat-Tree Networks	5
<i>German Maglione-Mathey (University of Castilla-La Mancha, Spain), Jesus Escudero-Sahuquillo (University of Castilla-La Mancha, Spain), Pedro Javier Garcia (University of Castilla-La Mancha, Spain), Francisco J. Quiles (University of Castilla-La Mancha, Spain), and Jose Duato (Universitat Politècnica de València, Spain)</i>	
High-Quality Fault-Resiliency in Fat-Tree Networks (Extended Abstract)	9
<i>John Gliksberg (UVSQ, Atos, UCLM), Antoine Capra (Atos), Alexandre Louvet (Atos), Pedro Javier García (UCLM), and Devan Sohier (UVSQ)</i>	

Session B: NOCs

Versal Network-on-Chip (NoC)	13
<i>Ian Swarbrick (Xilinx Inc.), Dinesh Gaitonde (Xilinx Inc.), Sagheer Ahmad (Xilinx Inc.), Bala Jayadev (Xilinx Inc.), Jeff Cuppett (Xilinx Inc.), Abbas Morshed (Xilinx Inc.), Brian Gaide (Xilinx Inc.), and Ygal Arbel (Xilinx Inc.)</i>	

Compute Express Link, Stephen Van Doren (Intel)	18
High Capacity On-Package Physical Link Considerations	19
<i>Greg Taylor (zGlue, Inc), Ramin Farjadrad (Aquantia), and Bapiraju Vinnakota (Netronome)</i>	

Session C: Links and FPGAs

Enabling Standalone FPGA Computing	23
<i>Joshua Lant (University of Manchester), Javier Navaridas (University of Manchester), Andrew Attwood (STFC), Mikel Lujan (University of Manchester), and John Goodacre (University of Manchester)</i>	
A Bunch of Wires (BoW) Interface for Inter-Chiplet Communication	27
<i>Ramin Farjadrad (Aquantia Corp) and Bapiraju Vinnakota (Netronome)</i>	
Demonstration of a Single-Lane 80 Gbps PAM-4 Full-Duplex Serial Link	32
<i>Sandeep Goyal (IIT BOMBAY), Pranav Agarwal (IIT BOMBAY), and Shalabh Gupta (IIT BOMBAY)</i>	

Session D: MPI and HPC

Improved MPI Multi-Threaded Performance using OFI Scalable Endpoints	36
<i>Aravind Gopalakrishnan (Intel Corporation, USA), Matias A. Cabral (Intel Corporation, USA), James P. Erwin (Intel Corporation, USA), and Ravindra Babu Ganapathi (Intel Corporation, USA)</i>	
Designing Scalable and High-Performance MPI Libraries on Amazon Elastic Fabric Adapter	40
<i>Sourav Chakraborty (The Ohio State University), Shulei Xu (The Ohio State University), Hari Subramoni (The Ohio State University), and Dhableswar Panda (The Ohio State University)</i>	
A Study of Network Congestion in Two Supercomputing High-Speed Interconnects	45
<i>Saurabh Jha (University of Illinois at Urbana-Champaign), Archit Patke (University of Illinois at Urbana-Champaign), Jim Brandt (Sandia National Lab), Ann Gentile (Sandia National Lab), Mike Showerman (NCSA), Eric Roman (NERSC), Zbigniew T. Kalbarczyk (UIUC), Bill Kramer (NCSA), and Ravishankar K. Iyer (UIUC)</i>	

Session E: Efficient Network Design and Use

Communication Profiling and Characterization of Deep Learning Workloads on Clusters with High-Performance Interconnects	49
<i>Ammar Ahmad Awan (The Ohio State University), Arpan Jain (The Ohio State University), Ching-Hsiang Chu (The Ohio State University), Hari Subramoni (The Ohio State University), and Dhableswar K. (DK) Panda (The Ohio State University)</i>	
Lightweight, Packet-Centric Monitoring of Network Traffic and Congestion Implemented in P4	54
<i>Philip Taffet (Rice University) and John Mellor-Crummey (Rice University)</i>	

OmniXtend: Direct to Caches Over Commodity Fabric	59
<i>Marjan Radi (Western Digital Research), Wesley W. Terpstra (SiFive, Inc), Paul Loewenstein (Western Digital Research), and Dejan Vucinic (Western Digital Research)</i>	
Latency Critical Operation in Network Processors	63
<i>Sourav Roy (NXP Semiconductors), Arvind Kaushik (NXP Semiconductors), Rajkumar Agrawal (NXP Semiconductors), Joseph Gergen (NXP Semiconductors), Wim Rouwet (NXP Semiconductors), and John Arends (NXP Semiconductors)</i>	
Author Index	69