Hot Interconnects 27

Ryan E. Grant, *Center for Computational Research, Sandia National Laboratories, Albuquerque, NM, 87185-1110, USA*

Manjunath Gorentla Venkata, *NVIDIA Corporation, Sunnyvale, CA, 94085, USA*

Welcome to the *IEEE Micro* Special Issue on High Performance Interconnects. In this issue, you will find the best papers on the hottest, most cutting edge interconnects design occurring in industry and academia from this year's IEEE Symposium on High Performance Interconnects (Hot Interconnects). Like our sister conference, Hot Chips, Hot Interconnects is a choice venue for revealing the latest and greatest advances in hardware. This year at Hot Interconnects was different than past events as it was virtual due to the COVID-19 pandemic. Hot Interconnects attendance was excellent, welcoming approximately 1000 attendees.

This year at Hot Interconnects we were pleased to welcome talks from network leaders from NVIDIA, AWS, Ayar Labs, and GigaIO. NVIDIA's CTO Michael Kagan laid out a case arguing that "Data center is a computer." Brian Barrett talked about the challenges of designing network technologies for AWS. Scott Taylor from GigaIO and Mark Wade from Ayar Labs introduced new network technologies.

OUR 2020 CONFERENCE FEATURED SEVERAL RIGOROUSLY PEER REVIEWED PAPERS, WITH AT LEAST FOUR REVIEWS PER SUBMISSION. THESE PAPERS SERVED AS THE BASIS FOR THE HIGH-QUALITY TALKS AT THIS YEAR'S HOT INTERCONNECTS.

A major theme of the submissions at this year's Hot Interconnects was SmartNICs. Following the theme, a panel lead by Scott Schweitzer from Xilinx had a very lively discussion on what constitutes a SmartNIC, how the implementation of SmartNIC effects the use cases, and the future vision of the SmartNIC's landscape. The experts in the discussion included leaders from NVIDIA, Xilinx, Broadcom, Fungible, and Pensando.

Our 2020 conference featured several rigorously peer reviewed papers, with at least four reviews per submission. These papers served as the basis for the high-quality talks at this year's Hot Interconnects. The top-rated papers from Hot Interconnects were invited to prepare new versions of their articles for inclusion in this special issue. These then went through an additional round of peer review followed by a rebuttal and revision phase.

The first article in this issue, "Configurable Network Protocol Accelerator (COPA)" by Krishnan *et al.*, describes the COPA framework, which enables programming field-programmable gate arrays (FPGAs) and use them for accelerating network protocols and scientific workloads. The article describes the hardware capabilities of FPGAs to reduce, transform, and process the network packets inline. Then, it describes integration into low-level middle layer called OpenFabric Interface. So, the functionality can be exposed to user in a hardware agnostic way. The value of the framework is shown by evaluating with workloads relevant to the scientific computing.

The article "Distributed Deep Learning With GPU-FPGA Heterogeneous Computing" by Tanaka *et al.* accelerates the distributed deep learning workloads by addressing the communication bottlenecks between GPUs. To address these bottlenecks, the paper leverages FPGA-based NICs to aggregate the data and execute the allreduce collective communication on the NICs. Besides that, the paper details the hardware and software ecosystem and demonstrating the value of the approach with the deep learning workloads.

In "PCI Express 6.0 Specification: A Low-Latency, High-Bandwidth, High-Reliability, and Cost-Effective Interconnect With 64.0 GT/s PAM-4 Signaling," Das Sharma details how PCI Express 6.0 will adopt pulse-amplitude modulation 4 (PAM-4) signaling to double the PCI Express 5.0 performance. The article details the error correction mechanisms required to balance the high bit error ratio of PAM-4 signaling. Additionally, the author also describes the techniques to achieve low latency, high bandwidth, and high reliability.

0272-1732 © 2021 IEEE
Digital Object Identifier 10.1109/MM.2020.3046118
Date of current version 26 January 2021.

The article "The Open Domain-Specific Architecture" by Vinnakota *et al.* addresses the problem of interconnecting chiplets on high-performance silicon. The article describes the open interfaces for both physical and logical die-to-die chiplet interfaces developed as part of the Open Compute Project efforts. In addition to an introduction to the concerns of chiplet interconnects, the authors describe the performance benefits to the proposed architecture using CCIX.

Olmedilla *et al.* provide a congestion isolation mechanism for lossy networks in "DVL-Lossy: Isolating Congesting Flows to Optimize Packet Dropping in Lossy Data-Center Networks." DVL-Lossy works by identifying and isolating congesting flows into special queues similar in design to dynamic virtual lanes that can be used in lossless networks. The goal of this work is to demonstrate that isolating congesting flows from noncongesting flows in different queues significantly reduces packet drops and increases network utilization and performance.

The article "Reliable and Time-Efficient Virtualized Function Placement" by Ben Haim and Rottenstreich explores the tradeoffs between reliability and time efficiency in networks. The authors identify three key factors that need to be taken into account for virtualized network functions, namely the number of chains, the length of said chains, and whether a given system allows function redundancy. They explore several optimizations in multiple scenarios, including those where multiple concerns intersect, finding a good compromise between competing factors when required.

The article "Chiplet Communication Link: Bunch of Wires (BoW)" by Ardalan *et al.* provides a chiplet die-to-die solution that offers a range of performance depending on multiple different levels of design/packaging. As the Bunch-of-Wires (BoW) interface is meant to enable heterogeneous integration, this article provides design and performance results from several major companies experiences with BoW including authors from Broadcom, Marvell, Ayar Labs, Keysight, and Apex.

WE HOPE THAT YOU WILL ENJOY THE BEST OF THIS YEAR'S HOT INTERCONNECTS AND JOIN US FOR HOT INTERCONNECTS 28 IN AUGUST 2021 AT QCT IN SAN JOSE, CA, USA.

In addition to the best of Hot Interconnects 27 in this issue, several other papers and content from our invited talks and keynotes are available in IEEE Xplore Digital Library at http://ieeexplore.ieee.org. We hope that you will enjoy the best of this year's Hot Interconnects and join us for Hot Interconnects 28 in August 2021 at QCT in San Jose, CA, USA. Participating in person is the only way to access some of the best interactive portions of the conference including our popular technical panel sessions and talks by industry leaders. In the day preceding the conference, we offer an array of technical tutorials to keep attendees up to date on the latest advances in networking and give them an opportunity to have hands on learning experiences facilitated by leaders from industry and academia. Further details can be found at http://www.hoti.org.

RYAN E. GRANT is a Principal Member of Technical Staff with the Center for Computational Research, Sandia National Laboratories, Albuquerque, NM, USA. He is a Senior Member of the IEEE. Contact him at regrant@sandia.gov.

MANJUNATH GORENTLA VENKATA is an HPC Software Architect with NVIDIA's Networking Unit, Sunnyvale, CA, USA. Contact him at manjunath@nvidia.com.