2nd North American Workshop on Silicon Photonics for High Performance Computing
May 29 – 30, 2019; Estes Park (Rocky Mountains), CO

Silicon-Photonic Networks for GPUs

Ajay Joshi
Boston University
joshi@bu.edu

Abstract

Silicon-photonic networks have been widely explored for inter-chip and intra-chip communication in manycore systems. However, there is limited work on the design of silicon-photonic networks for GPUs. In this talk, I'll present the limits and opportunities for using silicon-photonic networks for intra-GPU and inter-GPU communication. In particular, for the intra-GPU communication, I'll present our prior work on the design of a hybrid silicon-photonic network for GPUs. This hybrid network improves performance as well as energy efficiency. On the inter-GPU communication front, I'll present our ongoing work on jointly designing the GPU memory hierarchy and silicon-photonic network architecture for gaining high-performance energy-efficient execution of emerging workloads in the area of machine learning and graph processing.

Bio

Ajay Joshi is currently an Associate Professor in the ECE Department at Boston University. He received PhD from Georgia Tech in 2006 and then worked as a postdoc at MIT before joining Boston University in 2009. He worked as a Visiting Research Scientist at Google Inc. in 2017-18. He is currently the Associate Editor for IEEE Transactions on VLSI Systems. He has served on the technical program committee or external review committee of numerous conferences including DAC, DATE, NOCS, ISCA, MICRO, HPCA, ICCAD, ISQED and VLSI Design. He was a recipient of the Best Paper Award at AsiaCCS in 2018, Boston University ECE Award for Excellence in Teaching in 2014, NSF CAREER Award in 2012, and IEEE Micro Top Picks from Hot Interconnects in 2009. His current research interests include hardware security, silicon-photonic network architectures, hardware accelerators, and cross-layer optimization.