Performance evaluation of next-generation data center and HPC networks with co-packaged optics

Pavlos Maniotis

IBM Thomas J. Watson Research Center, Yorktown Heights, New York, 10598, USA
ppmaniotis@ibm.com

ABSTRACT

The increased escape bandwidth offered by co-packaged optics can enable switches with speeds of 51.2 Tb/s and beyond. From a network architecture perspective, there are two key advantages: (a) the implementation of large-scale topologies with significantly higher bisection bandwidth, and (b) the substantial reduction of the required number of switches, which can mitigate the administrative/management overhead. From a network operation perspective, both improved network locality and faster operation can be achieved since the higher-radix switches can reduce the impact of network contention; applications can be placed under fewer leaf switches, which reduces the number of packets that cross the spine switches in a typical leaf-spine topology. This presentation provides a brief overview of the recent activities realized within the framework of the MOTION research project (Multi-wavelength Optical Transceivers Integrated On Node), and discusses on the performance improvements that can be achieved by using co-packaged optics in next-generation data center and high-performance computing networks. The proposed concepts are evaluated via discrete-event simulations: first, virtual-machine traces are used to evaluate the network locality properties of the system, and, secondly, the performance improvements are quantified by means of network simulations with an OMNEST-based simulator.

Keywords: Network simulation, Network optimization, Co-packaged optics, Optical interconnects, Fat-tree networks.

RELATED MATERIAL

Previous papers with detailed results:

- P. Maniotis, L. Schares, D. M. Kuchta and B. Karacali, "Toward higher-radix switches with co-packaged optics for improved network locality in data center and HPC networks [Invited]," in Journal of Optical Communications and Networking, vol. 14, no. 6, pp. C1-C10, June 2022. https://ieeexplore.ieee.org/document/9729092
- P. Maniotis, N. Dupuis, L. Schares, B. G. Lee and D. M. Kuchta, "Intra-node High-performance Computing Network Architecture with Fast Optical Switch Fabrics," 2022 27th OptoElectronics and Communications Conference (OECC) and 2022 International Conference on Photonics in Switching and Computing (PSC), 2022, pp. 1-4. https://ieeexplore.ieee.org/document/9850165
- P. Maniotis, L. Schares, D. M. Kuchta and B. Karacali, "Improving Data Center Network Locality w/ Co-packaged Optics," 2021 European Conference on Optical Communication (ECOC), 2021, pp. 1-4. https://ieeexplore.ieee.org/document/9606112
- P. Maniotis, L. Schares, M. A. Taubenblatt and D. M. Kuchta, “Co-packaged optics for HPC and data center networks”, SPIE Photonics West, 6-11 March 2021, San Francisco, USA. https://www.spiedigitallibrary.org/conference-proceedings-of-spie/11692/1169205/Co-packaged-optics-for-HPC-and-data-center-networks/10.1117/12.2579066.full
- P. Maniotis, L. Schares, B. G. Lee, M. A. Taubenblatt and D. M. Kuchta, "Toward lower-diameter large-scale HPC and data center networks with co-packaged optics," in Journal of Optical Communications and Networking, vol. 13, no. 1, pp. A67-A77, January 2021, doi: 10.1364/JOCN.402676. https://ieeexplore.ieee.org/document/9275289
- P. Maniotis, L. Schares, B. G. Lee, M. A. Taubenblatt and D. M. Kuchta, "Scaling HPC Networks with Co-Packaged Optics," 2020 Optical Fiber Communications Conference and Exhibition (OFC), 2020, pp. 1-3. https://ieeexplore.ieee.org/document/9083409