Low Power Analog Coherent Links for Next-Generation Datacenters

Clint L. Schow
Dept. of Electrical and Computer Engineering, the University of California Santa Barbara, Santa Barbara, CA 93106, USA
schow@ece.ucsb.edu

Abstract: Ever-increasing bandwidth demand in datacenter networks makes a move to coherent links seem inevitable. An "analog coherent" approach using optical phase locked loops can enable low-power consumption, expanded link budgets, low-latency, and future bandwidth scalability.
OCIS codes: (200.0200) Optics in Computing; (060.2330) Fiber optics communications; (060.4265) Networks, wavelength routing

1. Introduction

Datacenters play critical roles in our everyday lives, our industries, and our societies. Datacenter traffic is expected to continue growing at a compound annual growth rate of 25% with server-to-server transfers accounting for >70% of the total [1]. Maximizing the bandwidth and efficiency of intra-datacenter communications is therefore key to improving overall datacenter productivity and efficiency. The INTREPID project, part of the ARPA-E ENLITENED program, is focused on developing a technology platform to integrate efficient high-speed photonic interfaces directly into switch IC packages [2]. Switch chips with co-packaged optical interfaces have the potential to be scaled to higher radix by eliminating the power-hungry electrical I/O currently required for connections to remote optical transceiver modules and applying the power savings to expand the core digital switch. The workhorse topology for cloud-scale datacenters is the fat-tree architecture, and larger radix switches can flatten these networks by removing layers of hierarchy, as illustrated in Figure 1. Flatter networks offer savings in power and cost with lower latency and improved bandwidth. Above the first switching layer in the fat-tree, often referred to as the top of rack (ToR), the links utilize single mode (SM) fiber due to its substantial advantages in operational management, cost, and support for future bandwidth scaling through wavelength division multiplexing (WDM). The longest of these SM intra-datacenter links are on the order of 2km.

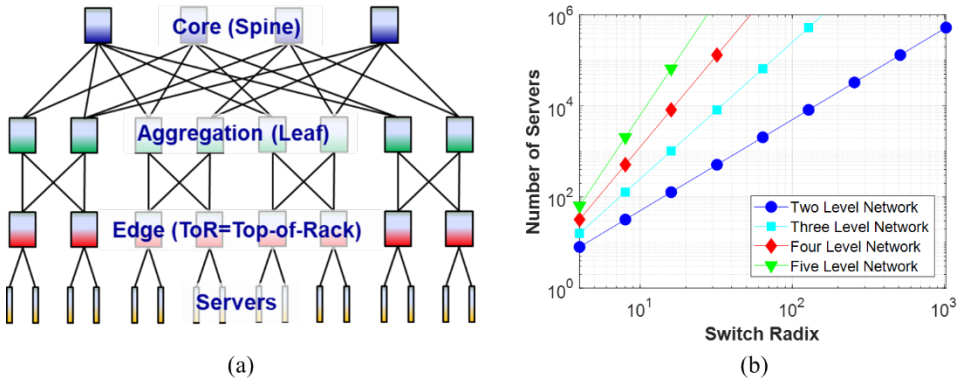

Figure 1. (a) Simple illustration of a fat-tree network, and (b) scaling properties: number of connected servers as a function of switch radix.

2. Analog Coherent Links for the Datacenter

For interconnects above the ToR tier, there is an opportunity for low-cost, low-power coherent WDM photonic interconnects purpose-built for the longer fabric links required in the datacenter. Tailoring the links to datacenter requirements requires optimization for a different set of metrics compared to current long-haul and metro coherent technology, specifically: 1) low power consumption, 2) expanded link budgets, 3) low cost, and 4) low latency. Moving datacenter interconnects from direct to coherent detection also enables future scalability to higher data rates through higher-order modulation formats and polarization modulation. For datacenters, the significantly larger link budget is a key advantage for coherent links. The additional link budget can potentially be applied to reducing link power by reducing the required optical power from the transmitter source laser or lowering cost by accommodating relaxed alignment tolerances and/or device loss specifications. Perhaps even more importantly, the higher sensitivity offered by coherent detection can be used to enable novel network architectures that incorporate all-optical routing/switching. Improvements in receiver sensitivity on the order of 20 dB are possible [3]. This key feature allows the coherent links to accommodate the loss of an AWGR (arrayed waveguide grating router) or active photonic switching layer in the optical network without requiring complex and costly integrated optical gain in such

components. Furthermore, the high selectivity offered by coherent reception significantly reduces the optical crosstalk requirements between channels for photonic routing/switching devices.

Coherent links for the datacenter can be drastically lower in power and complexity compared to current digital coherent technology that relies heavily on digital signal processing (DSP) to compensate for chromatic dispersion (CD), polarization mode dispersion (PMD), and nonlinear effects in DWDM links. In contrast, for < 2km links operating in the O-band near the zero-dispersion wavelength for standard single-mode fiber (1264-1338 nm), CD and PMD do not need to be compensated as they contribute negligible performance penalties. Furthermore, to eliminate the need for DSP-based carrier recovery, optical phase locked loops can be implemented that lock and track the phase and frequency of the receiver local oscillator (LO) to the incoming signal. The viability of highly-integrated OPLLs to enable robust and high-performance has been shown with previous demonstrations of analog coherent receivers operating "error-free" (< 10^{-12} BER up to 35 Gb/s) at 1550 nm [4].

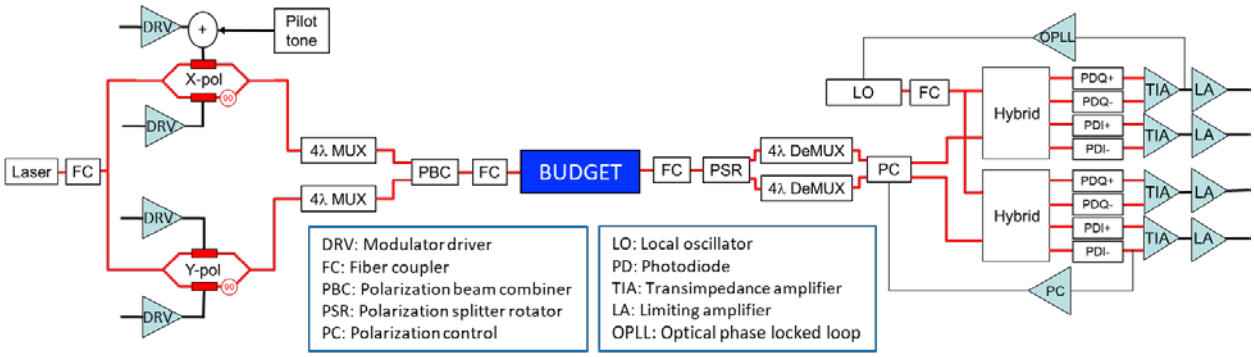

Figure 2. High-level schematic of a 200 Gb/s analog coherent link showing key functional blocks.

A block diagram of a 200 Gb/s analog coherent link is presented in Figure 2. The base symbol rate is 50 Gbd/s, the modulation format is QPSK, and polarization multiplexing is employed to maximize the bandwidth per fiber. The schematic is for a single wavelength, although four wavelengths are envisioned for 800 Gb/s/fiber with further scaling possible. Implementations in both InP and SiP photonic integrated circuit (PIC) technology are being pursued with a focus on maximizing the wall-plug power efficiency through photonic device and circuit co-design, choice of modulation format (QPSK), and close integration of electronics and photonics to minimize loop delays and maximize noise tolerance. The choice of QPSK supports low-power operation because after the in-phase (I) and quadrature phase (Q) signal components are separated by the 90° optical hybrid, simple and efficient limiting receivers can be used as only a binary decision is required. Therefore, linearity is not a concern and ADC followed by DSP is not required to achieve low native bit error rates. A preliminary analysis indicates that link budgets of 13 dB can be achieved while keeping wall-plug link efficiencies better than 10 pJ/bit.

3. Conclusion

Analog coherent links offer a promising path to increasing the bandwidth of datacenter interconnects while minimizing power consumption and enabling future network architectures incorporating all-optical routing/switching. When integrated directly into first-level chip packages, switches can be built with larger port count to further improve data center performance and efficiency by flattening the network through eliminating layers of hierarchy.

4. Acknowledgments

The author greatly appreciates the technical contributions and rewarding collaborations with colleagues at UCSB (James Buckwalter, Larry Coldren, Jonathan Klamkin, Adel Saleh, Shamsul Arafin, Shireesh Bhat, Sarvagya Dwivedi, Fabrizio Gambini, Sergio Pinna, Hector Andrade, Takako Hirokawa, Aaron Maharry, Thomas Meissner, Stephen Misak, Luis Valenzuela, Yujie Xia) and Facebook (Katharine Schmidtke, Hans-Juergen Schmidtke, Ariel Hendel, Brian Taylor, Todd Hollmann, Jimmy Williams, Gilad Goldfarb, James Stewart).

Funding under the ARPA-E ENLITENED program is gratefully acknowledged and the authors would like to especially thank the ARPA-E team for management and guidance: Michael Haney, James Zahler, and Alan Liu. The information, data, or work presented herein was funded in part by the Advanced Research Projects Agency-Energy (ARPA-E), U.S. Department of Energy, under Award Number DE-AR0000848. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

5. References

[1] White Paper: "Cisco global cloud index: forecast and methodology, 2016–2021," available: https://www.cisco.com/c/en/us/solutions/collateral/service-provider/global-cloud-index-gci/white-paper-c11-738085.pdf, accessed 12/12/2018.

[2] C. L. Schow and K. E. Schmidtke, "*INTREPID:* Developing Power Efficient Analog Coherent Interconnects to Transform Data Center Networks," *OFC 2019*, San Diego, CA, Mar. 2019.

[3] J. K. Perin, A. Shastri and J. M. Kahn, "Design of Low-Power DSP-Free Coherent Receivers for Data Center Links," in *J. Lightw. Technol.*, vol. 35, no. 21, pp. 4650-4662, Nov., 2017.

[4] M. Lu *et al.*, "An Integrated 40 Gbit/s Optical Costas Receiver," *J. Lightw. Technol.*, vol. 31, no. 13, pp. 2244-2253, July, 2013.