Transceivers and Optical Engines for Computer and Datacenter Interconnects

Mitchell H. Fields, John Foley, Ron Kaneshiro, Larry McColloch, David Meadowcroft, Frederick W. Miller, Sami Nassar, Michael Robinson, Hui Xu
Avago Technologies, 350 W. Trimble Road, San Jose, CA 95131
Author email: mitch.h.fields@avagotech.com

Abstract: Technological and manufacturing advances are enabling optical solutions to meet demanding density, bandwidth, and cost metrics. We describe these solutions with respect to different system architectures and applications.

OCIS code: (200.4650) Optics in computing, optical interconnects

1. Introduction

Modern datacenters, high-performance computing (HPC) centers, and supercomputers demand dense cost-effective high-bandwidth interconnect solutions between various computing resources and network elements. High-volume low-cost manufacturing technologies are as critical as the fundamental electro-optics technologies in delivering these solutions. Here, we review Avago Technologies' building-block / optical-engine approach to address both single-channel and multi-channel transceiver solutions.

2. Single-channel interconnects

Single channel SFP and SFP+ transceivers are now the dominant form factors for multimode 1/10 GbE and 1/2/4G FC interconnects in the data center for servers, switches, routers, and storage devices. Traditional transceiver designs are based upon discrete optical subassemblies (OSAs) for the transmitter (TOSA) and receiver (ROSA). These OSAs are attached to the module PCB via electrical leads or flex circuits. In contrast to this method, Avago Technologies has adopted a "bulk-optics" approach to assembly of the transceivers that avoids manual handling of OSAs, thus mitigating potential risks of unintentional damage to the high-speed III/V devices (e.g., electrostatic, mechanical). An exploded view of the assembly is shown in Figure 1. With the PCBs in panel form, a fully automated process places the VCSEL and photodiode devices directly onto the transceiver PCB. A low-cost plastic lens assembly is attached over the active devices to provide the lensing, right-angle turn, and optical port functionality. Depanelization and manual handing of the SFPs occurs *after* the sensitive III/V devices are protected from potential external damage. The PCBs are then used as engines in SFP assemblies for 1 – 10G applications. The volume enjoyed by this building-block approach allows for low-cost manufacturing.

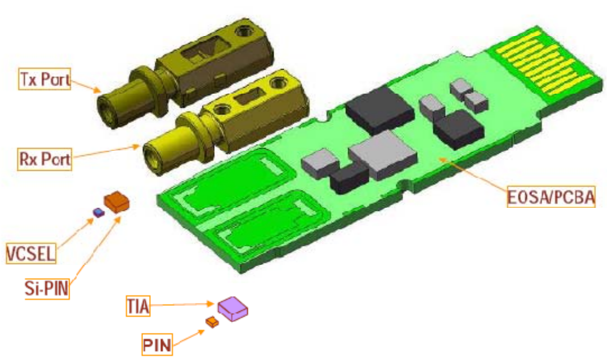

Figure 1: Bulk optics approach to transceiver assembly.

3. Parallel multi-channel interconnects

In datacenter and HPC applications, the performance of the system is oftentimes determined by the bisection bandwidth of the fabric interconnect. This is true for the fabric between clustered nodes as well as for the internal fabric used by multi-chassis switches and routers. Parallel-optic modules are the ideal solution for this connectivity problem, though the cost can sometimes be prohibitive, thus forcing a suboptimal architecture that relies upon a copper interconnect. In situations where the bandwidth-density can be accommodated by front-panel mounted transceivers, QSFP and CXP form factors can provide a solution. However, there are applications where signal-integrity, thermal and density requirements preclude the use of front-panel mount optics. To address these cost and application concerns, a significant investment is required to realize a paradigm shift in the design, manufacture, and test of parallel-optic modules. As a datacenter focused optical supplier, Avago Technologies has developed a novel embedded parallel-optic module designed for high-volume manufacturing in panel form. This module is available as a discrete component for customers and as an engine for Avago-designed integrated parallel-optics assembles. The 12-channel module provides over 120 Gb/s of connectivity in less than 80 mm^2. The module can be placed near the host ASIC to mitigate signal-integrity concerns. A unique low-cost optical connector affixes to the top of the module to allow for tight packing of multiple devices for dense applications. Custom designed control ICs provide features such as diagnostic test and monitoring as well as flexibility for signal-integrity enhancement.

4. Summary

We have described how transceivers based upon a building-block / optical-engine approach can successfully address the interconnect requirements of the datacenter. Furthermore, the approach leads to cost-effective solutions that are commercially viable.