Trends and Future Directions in Optical Interconnects for Datacenter and Computer Applications

Katharine Schmidtke and Frank Flens
Finisar Corporation, 1389 Moffett Park Drive, Sunnyvale, CA 94089
katharine.schmidtke@finisar.com

Abstract: The underlying technology trends that have enabled optics to displace copper in telecommunications, datacommunications, and datacenters are discussed. Future directions are outlined that will drive optical interconnects further into the fabric of computer systems.

OCIS codes: (230.0250) Optoelectronics; (200.4650) Optical Interconnects

1. Trends in optical interconnects over the past decade

The optical transceiver market has seen unprecedented growth over the past decade, and transceivers are now ubiquitous in telecommunication and datacommunication deployments worldwide. The primary driver behind the use of optical links has been bandwidth-reach and the inability of copper cable to support increases in bandwidth over the required link lengths. This was true in telecommunication applications in the1980s, and then in datacommunications in the 1990's and is now the case in datacenters where optical interconnects rival copper cables as high bandwidth interconnects.

Optical transceivers have revolutionized communication systems in a number of ways over the past decade. As bandwidth requirements have increased, so has the pressure to maintain small footprints and low power budgets. The graph below shows how the package size and power dissipation has shrunk over the past decade through various generations of optical modules. Optoelectronic integration and package miniaturization have been crucial in achieving these reductions producing an increase in the bandwidth density of more than two orders of magnitude over the past ten years.

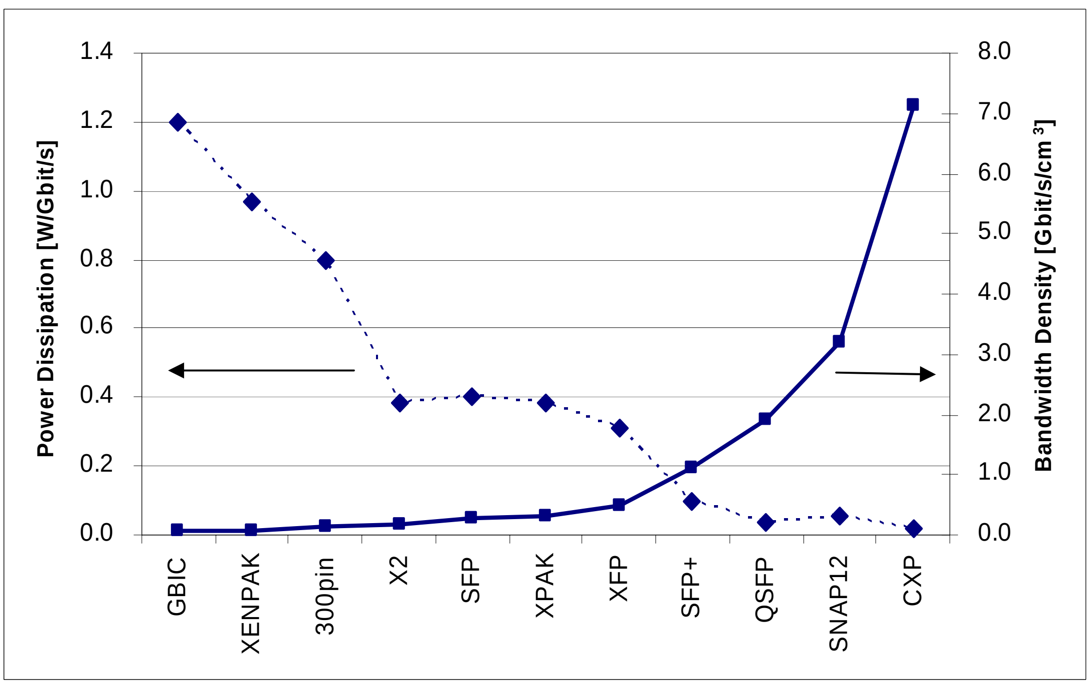

Fig. 1. Graph illustrating the miniaturization and power efficiency trends in optical transceivers over the past decade.

Standardization and pluggability have also played an important role in shaping the communications industry. Standardization has many obvious benefits in a commercial environment, driving volume on a standard manufacturing platform. Pluggabilty is an innovation that allows a single basic line-card or host board to support any link length, from meters to kilometers, simply by plugging in the appropriate type of optical transceiver from a given form-factor. Front panel pluggability also allows links to be installed as required in a "pay as you grow" implementation.

2. Optical interconnects in the datacenter – active optical cables

Optical interconnects are extending their applicability into new areas beyond specialized trans-continental or cross-campus links. Recent deployments have used optical cables to connect between racks of equipment within the datacenter, and within high performance computer clusters. The world's fastest supercomputer, "Roadrunner" [1], was deployed in 2008 and uses optical interconnects exclusively for inter-rack interconnects. Optical interconnects are becoming standard practice for datacenter installations and within the last two years, active optical cables have begun to rival copper cables for high bandwidth links. InfiniBand is a communications link that is known for low latency and high bandwidth and is used in datacenters and high performance computers. The InfiniBand Trade Association was an early adopter of active optical cables, adding optics as an approved interface in 2007 [2].

Active optical cables mimic copper cables in terms of their ease of use, and offer a number of key advantages over copper such as bandwidth-reach. As bandwidth increases, the length of interconnect that copper cables can support decreases. At 10 Gbit/s signaling rate, copper supports less than 10 m. In addition, optical cables are much thinner. For example, a cable with four full-duplex links at 10 Gbit/s in optical ribbon fiber has a cross section of just 2 mm x 4 mm, compared to a similar link in copper cables which have diameters of more than 6 mm. The associated reduction in weight and cabling hardware allows significant saving on the installation infrastructure. Thinner cables also allow better air-flow through the system, reducing the power required to cool the datacenter. Further, the total power consumption of the link is lower using optics, even compared to passive copper cables, because there is no need for complicated error correction schemes or signal pre-emphasis.

There are two main types of active optical cables used in datacenters today with form factors known as QSFP [3] and CXP [4]. The two form-factors are illustrated in Fig. 2 below. QSFP provides four parallel lanes in each direction, and CXP has twelve parallel lanes in each direction. The optical link is capable of 12.5 Gbit/s per lane providing aggregate bandwidth of up to 150 Gbit/s full-duplex for proprietary applications. What is significant here is that the optical interface is contained within the plug end, creating an enclosed optical system. These are not transceivers with detachable optical fibers, they are integrated cables that look and feel very much like copper cables, but with improved performance and less bulk.

Fig. 2. Examples of active optical cables used in datacenter applications: QSFP (left) and CXP (right).

Contained within the active optical cables shown in Fig. 2, are bundles of parallel optical fibers arranged in a ribbon format. QSFP has four lanes in direction lighting up 8 fibers, and CXP with twelve lanes uses 24 fibers. Taking InfiniBand QDR datarate as a particular example, each separate lane runs at 10 Gbit/s corresponding to an aggregate bandwidth of 40 Gbit/s for the QSFP and 120 Gbit/s for the CXP form factor.

OTuP4.pdf

3. Future trends for the next decade

Bandwidth-reach will remain a driving force behind the adoption of optical interconnects over the next decade and will bring optics deeper into the fabric of computer systems [5]. Every product generation will continue to demand improvements in bandwidth while maintaining link length and small size. The InfiniBand Trade Association is already working on defining the next generation specification called EDR which will support double the bandwidth of today's QDR specification. At datarates of 25 Gbit/s copper cables, and even copper traces on host boards become extremely challenging [6]. Optical links are not trivial at these datarates but offer many overwhelming advantages over copper including bandwidth-reach, reduced power and immunity to electro-magnetic interference.

In addition to bandwidth-reach, a shift of focus is occurring towards bandwidth-density. This will be especially true for input/output (I/O) interconnects – all the way from cabling between racks down to interconnects between chips [7]. The ability to fabricate hundreds of parallel optical links on 250 micron spacing will go a long way to alleviating the next I/O bottleneck between computer boards and between chips. As process power increases, innovative technologies in photonic integration, packaging and automation will reduce the foot-print, increase the functionality and enable mass manufacturing of optical interconnects.

It is clear from the insatiable demand for bandwidth throughout the entire network infrastructure that the trend to miniaturization will continue long into the future.

4. References

[1] List of Top 500 Supercomputers as compiled by *Top 500.org* (June 2008).
[2] Alan Benner, "Annex A5: Pluggable Interfaces: CATx, Copper and Optical" White Paper (IBTA Electromechanical Working Group) (March 2007).
[3] The specification for QSFP+ "Copper and Optical Modules" is documented by the SFF committee and can be found at ftp://ftp.seagate.com/pub/sff/SFF-8436.PDF
[4] The specification for CXP "Mini Multilane Series" is documented by the SFF committee and can be found at ftp://ftp.seagate.com/pub/sff/SFF-8642.PDF
[5] Terry Morris, "Breaking free of electrical constraints," Appl. Phys.A **95,** 941-944 (2009)
[6] Dong Kam et al. "Is 25 Gb/s on-board signaling viable?" IEEE Transactions on Advanced Packaging, (2009)
[7] Ashok Krishnamoorthy, "Optical proximity communication with passively aligned silicon photonic chips," IEEE Journal of Quantum Electronics, Vol. 45**,** pp. 409-414 (2009)