Volume 08 Issue 02 Published, May 10, 2004 ISSN 1535-864X

Intel® Technology Journal

Optical Technologies and Applications

Optical Technologies for Enterprise Networks

Optical Technologies for Enterprise Networks

Pierre Herve, Intel Communications Group, Intel Corporation
Shlomo Ovadia, Intel Communications Group, Intel Corporation

Index words: optical interconnects, 10 Gb/s transceivers, packaging, enterprise networks

ABSTRACT

Optical networking technologies have been over the last two decades reshaping the entire telecom infrastructure networks around the world. As network bandwidth requirements increase, optical communication and networking technologies have been moving from their telecom origin into the enterprise. For example, today in data centers, all storage area networking is based on fiber interconnects with speeds ranging form 1 Gb/s to 10 Gb/s. As the transmission bandwidth requirements increase and the costs of the emerging optical technologies become more economical, the adoption and acceptance of these optical interconnects within enterprise networks will increase. This paper, which provides the framework for the different optical interconnect technologies in this special optical issue of the *Intel Technology Journal*, is organized as follows. First, a brief overview of the fiber optics interconnects technology evolution and its current application within the enterprise, is presented. Second, various interconnect evolution paths, such as, board-to-board, chip-to-chip, and on-chip interconnects, are discussed.

INTRODUCTION

The birth of optical communications occurred in the 1970's with two key technology breakthroughs. The first was the invention of the semiconductor laser in 1962 [1]. The laser generates a tightly focused beam of light at a single pure wavelength, a spot small enough to be connected to fiber optics. The second breakthrough happened in September 1970, when a glass fiber with an attenuation of less than 20 dB/km was developed [2, 3]. In the 1960's, glass-clad fibers had an attenuation of about 1 dB/m, which was sufficient for medical imaging applications, but was too high for telecommunications. With the development of optical fibers with an attenuation of 20 dB/km, the threshold to make fiber optics a viable technology for telecommunications was crossed. In 1977, AT&T installed the first optical fiber cables in Chicago [3]. The first field deployments of fiber communication systems used Multimode Fibers (MMFs) with lasers operating in the 850 nm wavelength band. These systems could transmit several kilometers with optical losses in the range of 2 to 3 dB/km. A second generation of lasers operating at 1310 nm enabled transmission in the "second window" of the optical fiber where the optical loss is about 0.5 dB/km in a Single-Mode-Fiber (SMF). In the 1980's, the telecom carriers started replacing all their MMFs operating at 850 nm. Another wavelength window around 1550 nm was developed where a standard SMF has its minimum optical loss of about 0.22 dB/km. The development of fiber-based telecommunication systems in the 1990's focused on increasing their transmission capacity. This was done first by increasing the signal modulation speed from 155 Mb/s to 622 Mb/s, to 2.5 Gps, and finally to 10 Gb/s, today's modulation speed. The total available bandwidth of standard optical fibers is enormous; it is about 20 THz. Since it is impossible for a single-wavelength laser to utilize this enormous bandwidth, multiple single-wavelength laser transmitters are typically multiplexed and transmitted on a single fiber. This scheme, which was developed in the mid 1990's, is called Wavelength-Division-Multiplexing (WDM) [4]. Dense WDM (DWDM) optical communication systems with more than 60 wavelengths, where each wavelength carries 40 Gb/s data, have been demonstrated [5]. Thus, the demonstrated total transmission capacity of an SMF is more than 2.5 Tb/s.

Today, MMFs operating at 850 nm are primarily used for short distances in the enterprise as the least expensive method. An SMF at the 1310 nm wavelength band is primarily used for medium distances ranging from 2 km to 40 km. For long-haul telecommunications, WDM systems operating in the 1550 nm wavelength band windows are deployed. From 850 nm to long wavelength and WDM, higher performance is being offered, but each one comes with a higher price tag. Nowadays, all the telecom infrastructure is fiber-based with the exception of the famous last miles to homes, which is still based on coaxial cables and copper-twisted pairs. Inside enterprise networks, fiber has been deployed since the early 1980's initially with supercomputers, and later in Local Area Networks (LANs) as well as more recently in Storage Area Networks (SANs). With continuously increasing demands for high-speed data, optical fibers and interconnects will continue to play an increasing role within the enterprise network.

In this paper, we first discuss the fundamentals of optical components for communication such as optical fibers, laser transmitters, and receivers. Second, we review the optical packaging trends for optical modules and optical transceivers. Then, after a review of the various optical interconnect topologies in the enterprise, we discuss their applications and cost trends. We then move on to discuss the next-generation optical interconnects dividing them into four main categories: box-to-box, board-to-board, chip-to-chip, and on-chip interconnects. Finally, the evolution of optical interconnects and technical challenges in enterprise networks are discussed.

THE FUNDAMENTALS OF OPTICAL COMPONENTS

A basic optical communication link consists of three key building blocks: optical fiber, light sources, and light detectors. We discuss each one in turn.

Optical Fibers

In 1966, Charles Kao and George Hockmam predicted that purified glass loss could be reduced to below 20 dB per kilometer, and they set up a world-wide race to beat this prediction. In September 1970, Robert Maurer, Donald Keck, and Peter Schultz of Corning succeeded in developing a glass fiber with attenuation less than 20 dB/km: this was the necessary threshold to make fiber optics a viable transmission technology. The silica-based optical fiber structure consists of a cladding layer with a lower refractive index than the fiber core it surrounds. This refractive index difference causes a total internal reflection, which guides the propagating light through the fiber core. There are many types of optical fibers with different size cores and cladding. Some optical fibers are not even glass-based such as Plastic Optical Fibers (POFs), which are made for short-distance communication. For telecommunications, the fiber is glass based with two main categories: SMF and MMF. SMFs typically have a core diameter of about 9 μm while MMFs typically have a core diameter ranging from 50 to 62.5 μm. Optical fibers have two primary types of impairment, optical attenuation and dispersion. The fiber optical attenuation, which is mainly caused by absorption and the intrinsic Rayleigh scattering, is a wavelength-dependent loss with optical losses as low as 0.2 dB/km around 1550 nm for conventional SMF (i.e., SMF-28[*]) [6].

The optical fiber is a dispersive waveguide. The dispersion results in Inter Symbol Interference (ISI) at the receiver. There are three primary types of fiber dispersions: modal dispersion, chromatic dispersion, and polarization-mode dispersion. The fiber modal dispersion depends on both the fiber core diameter and transmitted wavelengths. For a single-mode transmission, the step-index fiber core diameter (*D*) must satisfy the following condition [2]:

$$D < \left[\frac{2.405 \cdot \lambda}{\pi}\right] \cdot \left(n_1^2 - n_2^2\right)^{-1/2}$$

where λ is the transmitted wavelength and n_1 and n_2 are the refractive indices of fiber core and cladding layer, respectively. Consequently, for a single-mode operation at 850 nm wavelength, the fiber must have a core diameter of 5 μm. Since a conventional SMF has typically a core diameter of 9 μm, single-mode operation can be only supported for wavelengths in the 1310 nm wavelength band or longer.

The fiber chromatic dispersion is due to the wavelength-dependent refractive index with a zero-dispersion wavelength occurring at 1310 nm in conventional SMF [6]. At 1550 nm, the fiber dispersion is about 17 ps/nm/km for SMF-28. When short duration optical pulses are launched into the fiber, they tend to broaden since different wavelengths propagate at different group velocities, due to the spectral width of the emitter. Optical transmission systems operating at rates of 10 Gb/s or higher and distances above 40 km are sensitive to this phenomenon. There are other types of SMFs such as Dispersion Shifted Fibers (DSFs) where the zero dispersion occurs at 1550 nm.

Polarization-Mode Dispersion (PMD) is caused by small amounts of asymmetry and stress in the fiber core due to the manufacturing process and environmental changes such as temperature and strains. This fiber core asymmetry and stress leads to a polarization-dependent index of refraction and propagation constant, thus limiting the transmission distance of high speed ($\geq$ 10 Gb/s) over SMF in optical communication systems. Standard SMF has a PMD value of less than 0.1 ps/√km [6]. Special SMFs were developed to address this issue.

---

[*] All other trademarks are the property of their respective owners.

Optical fiber is never bare. The fiber is coated with a thin primary coating by the fiber manufacturer; then a cable manufacturer, not necessarily the fiber manufacturer, cables the fiber. There is a wide variety of cable construction. Simplex cable has a single fiber in the center while duplex cables contain two fibers. Composite cable incorporates both single-mode and multimode fiber. Hybrid cables incorporate mixed optical fiber and copper cable. In the enterprise, the MMF is housed in a cable with an orange colored jacket, and the SMF is housed in a yellow jacket cable.

Light Sources

The light source is often the most costly element of an optical communication system. It has the following key characteristics: (a) peak wavelength, at which the source emits most of its optical power, (b) spectral width, (c) output power, (d) threshold current, (e) light vs. current linearity, (f) and a spectral emission pattern. These characteristics are key to system performance.

There are two types of light sources in widespread use: the Laser Diode (LD) and the Light Emitting Diode (LEDs). All light emitters that convert electrical current into light are semiconductor based. They operate with the principle of the p-n semiconductor junction found in transistors. Historically, the first achievement of laser action in GaAs p-n junction was reported in 1962 by three groups [1-4]. Both LEDs and LDs use the same key materials: Gallium Aluminum Arsenide (GaAIAs) for short-wavelength devices and Indium Gallium Arsenide Phosphide (InGaAsP) for long-wavelength devices.

Semiconductor laser diode structures can be divided into the so-called edge-emitters, such as Fabry Perot (FP) and Distributed Feedback (DFB) lasers and vertical-emitters, such as Vertical Surface Emitting Lasers (VCSELs). When edge-emitters are used in optical fiber communication systems, they incorporate a rear facet photodiode to provide a means to monitor the laser output, as this output varies with temperature.

In today’s optical networks, binary digital modulation is typically used, namely on (i.e., light on) and off (no light) to transmit data. These semiconductor laser devices generate output light intensity which is proportional to the current applied to them, therefore making them suitable for modulation to transmit data. Speed and linearity are therefore two important characteristics.

Modulation schemes can be divided into two main categories, namely, a direct and an external modulation. In a direct modulation scheme, modulation of the input current to the semiconductor laser directly modulates its output optical signal since the output optical power is proportional to the drive current. In an external modulation scheme, the semiconductor laser is operating in a Continuous-Wave (CW) mode at a fixed operating point. An electrical drive signal is applied to an optical modulator, which is external to the laser. Consequently, the applied drive signal modulates the laser output light on and off without affecting the laser operation.

One important feature of the laser diode is its frequency chirp. The frequency of the output laser light changes dynamically in response to the changes in the modulation current. A typical DFB has a frequency chirp of about 100-MHz/mA. This spread of the wavelength interacts with the fiber dispersion. As previously mentioned, as the data rate is increased, this interaction limits the transmission distance of optical transmission systems due to the additional ISI generated at the receiver [1–4].

Optical back-reflection is one of key issues when coupling the output light from a laser source to a fiber. The optical back-reflection disturbs the standing wave in the laser cavity, increasing its noise floor, and thus making the laser unstable. One practical way to reduce the phenomenon of back-reflection is to place an isolator between the laser cavity and the fiber, which adds a significant additional cost to the laser [1, 4]. Temperature also affects the peak wavelength of the laser; threshold current also increases with temperature as slope efficiency decreases. For DWDM applications, which require very precise operating wavelengths, most of the current laser diode designs need to be cooled to within ± 0.3 °C.

As previously explained, the direct modulation of a laser diode has several limitations, including limited propagation distance due to the interaction between the laser frequency chirp and fiber dispersion. This is not an issue for enterprise networks which are short distance, but could be a serious limiting factor for telecommunications applications. To overcome this limitation, the laser diode is operated in a CW mode, and output light is externally modulated by an optical modulator. Intensity modulators can be divided into two main groups: Mach-Zehnder Interferometer (MZI) and Electro-Absorption (EA) modulators. In an MZI modulator, a single input waveguide is split into two optical waveguides by a 3 dB Y junction and then recombined by a second 3 dB Y junction into a single output. A Radio Frequency (RF) signal, which is applied to a pair of electrodes constructed along the waveguides, modulates the propagating optical beam. The modulator key parameters are its modulation bandwidth, linearity, and the required drive signal voltage for π phase shift. MZI modulators based on $LiNbO_3$ are high-performance modulators with a large form-factor (about 2.5 inches) that are not suitable for optical integration [4, 7]. EA modulators are based on a voltage-induced shift of the semiconductor bandgap so that the modulator becomes absorbing for the lasing wavelength. The advantages of an EA modulator is its low driving voltage, high-speed operation, and suitability for optical integration with InP-based laser diodes [8].

A tunable laser is a new type of laser where its main lasing longitudinal mode can be tuned over a wide range of wavelengths such as the C band (1510–1540 nm) of an Erbium-Doped Fiber Amplifier (EDFA), which is commonly used for DWDM systems [1–4]. The use of tunable lasers is driven by the potential cost savings in DWDM transport networks since a significantly reduced inventory of fixed-wavelength lasers could be maintained for a robust network operation. The technical challenges are to provide both broad wavelength tunability and excellent wavelength accuracy over the laser life. A broadly tunable External Cavity Laser (ECL) employing micromachined, thermally tuned silicon etalons has been designed to achieve these goals.

Light Detectors

Light detectors convert an optical signal to an electrical signal. The most common light detector is a photodiode. It operates on the principle of the p-n junction. There are two main categories of photodetectors: a p-i-n (positive, intrinsic, negative) photodiode and an Avalanche Photodiode (APD), which are typically made of InGaAs or germanium. The key parameters for photodiodes are (a) capacitance, (b) response time, (c) linearity, (d) noise, and (e) responsivity. The theoretical responsivity is 1.05 A/W at a wavelength of 1310 nm. Commercial photodiodes have responsivity around 0.8 to 0.9 A/W at the same wavelength [1-4]. The dark photo-current is a small current that flows through the photo-detector even though no light is present because of the intrinsic resistance of the photo-detector and the applied reverse voltage. It is temperature sensitive and contributes to noise. Since the output electrical current of a photodiode is typically in the range of μA, a Transimpedance Amplifier (TIA) is needed to amplify the electric current to a few mA [2–4].

APDs provide much more gain than the pin photodiodes, but they are much more expensive and require a high voltage power to supply their operation [2]. APDs are also more temperature sensitive than pin photodiodes.

PACKAGING: PACKAGING: OPTICAL SUB-ASSEMBLY (OSA) AND OPTICAL TRANSCEIVERS

As previously described, laser diodes and photodiodes are semiconductor devices. To enable the reliable operation of these devices, an optical package is required. In general, there are many discrete optical and electronic components, which are based on different technologies that must be optically aligned and integrated within the optical package. Optical packaging of laser diodes and photodiodes is the primary cost driver. These packages are sometimes called Optical Sub-Assemblies (OSAs). The Transmitter OSA package is called a TOSA and the Receiver OSA package is called a ROSA.

Figure 1 shows, for example, a three-dimensional schematic view of a DFB laser diode mounted on a Thermo-Electric Cooler (TEC) inside a hermetically sealed 14-pin butterfly package with an SMF pigtail [9]. Most of the telecom-grade laser diodes are available in the so-called TO can or butterfly packages. The standard butterfly package is a stable and high-performance package, but it has a relatively large form-factor and it is costly to manufacture. These packages are typically used for applications where cooling is required using a TEC [4].

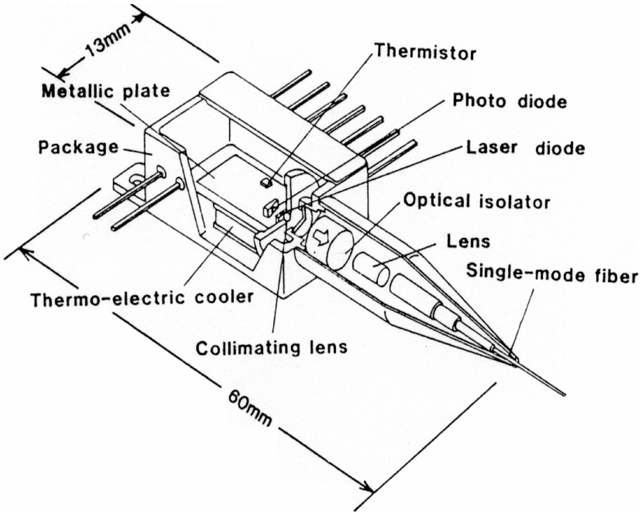

Figure 1: Three-dimensional view of a DFB laser diode configuration with single-mode fiber pigtail (after Ref. [8] (© 1990 IEEE))

The TEC requires a large amount of power to regulate the temperature of a laser inside the package. This type of optical packaging was used for the early 10 Gb/s modules. More recently, tunable 10 Gb/s lasers are using a similar butterfly optical package. The butterfly package design uses a coaxial interface for passing broadband data into the package, which requires the use of a coaxial interface to the host Printed Circuit Board (PCB). Although coaxial cables and connectors have been reduced in size, they still consume valuable real estate in the optical transceiver.

The evolution of optical module packages is toward smaller footprint packages. If relatively easy for receivers, the trend toward smaller packages is particularly challenging for laser transmitter modules due to the power and thermal dissipation constraints. Figure 2 shows the evolution of 10 Gb/s optical module packaging technology. To operate with high-performance, uncooled designs must be implemented with more advanced control systems that can adjust the laser and driver parameters over temperature. The smaller packages utilize a coplanar approach to the broadband interface, which more closely resembles a surface-mount component and enables much smaller RF interfaces.

TO-can-based designs, which have been used extensively in lower data rate telecom and datacom systems up to 2 Gb/s as well as CD players and other high-volume consumer applications, are now maturing to support high-performance 10 Gb/s optical links. Leveraging the fact that these packages are already produced in high volume will further reduce the cost of the 10 Gb/s optical modules in optical transceiver designs.

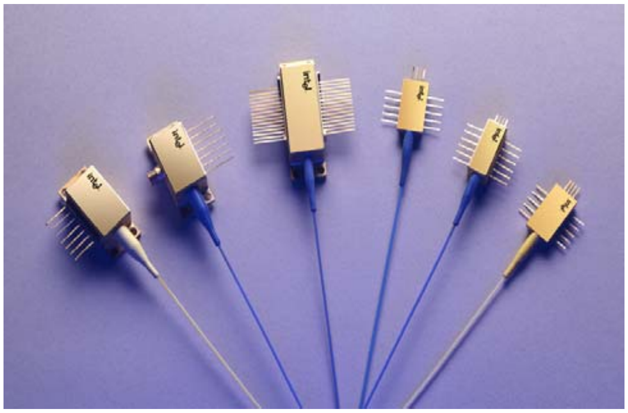

Figure 2: Trends in 10 Gb/s optical transmitter packaging technology. To decrease size and power dissipation, the trends are from cooled to uncooled packages, from coaxial to planar RF interfaces, and from pigtailed to pluggable optical interfaces.

The current TOSA/ROSA package form factors are trending toward smaller packages, and it will not end at the 10 Gb/s TO-can implementation. These TOSA/ROSA form factors are still too large and too expensive to compete in the market segment where today copper interconnects dominate. Leadframe-type packages could be an attractive choice for high-speed optical modules since similar packages are already in use in the semiconductor industry [10]. Using insert molded or pre-molded thermoplastic housing, different optical components can be passively aligned in a fully automated manufacturing process. For example, integrated modules with VCSEL and photodiodes in leadframe packages have been developed for the automobile industry, but are limited today in the 20 Mb/s bit rate. Additional research and development is needed to define the packaging specifications for optoelectronic modules based on size and cost.

Optical Transceivers

For telecommunication applications, the optical transmitter and receiver modules are usually packaged into a single package called an optical transceiver. Figure 3 shows an example of different transceivers and Figure 4 shows an example of the printed circuit board of a transceiver. There are several form factors for this optical transceiver depending on their operating speed and applications. The industry worked on a Multi-Source Agreement (MSA) document to define the properties of the optical transceivers in terms of their mechanical, optical, and electrical specifications. Optical transponders operating at 10 Gb/s, based on MSA, have been in the market since circa 2000, beginning with the 300-pin MSA, followed by XENPAK, XPAK, X2, and XFP. Table 1 summarizes the key MSA specifications for the different form-factor 10 Gb/s optical transceivers and their release dates.

Which are the Most Popular Form-Factor Transceivers in the Enterprise?

For the 1/2/4 Gb/s transceivers, the Small Form-Factors (SFFs) and the small Form-Factor Pluggables (SFPs) are the most recently developed and the ones that are finding new sockets into systems. It should be noted, however, that the older GBIC form factors for 1 Gb/s Ethernet (GbE), despite no new development, is still shipping in large volumes due to the large installed base of this design. The SFF transceiver is used in a Network Interface Card (NIC) for the LAN or in the Host Bus Adaptor (HBA) in SANs. The SFP transceiver is typically used for enterprise switches such as Ethernet or Fiber-Channel (FC) switches. In these high-capacity switches, switching is done by electrical ICs while the optical transceivers provide optical-to-electrical electrical (O-E) or electrical-to-optical (E-O) conversion.

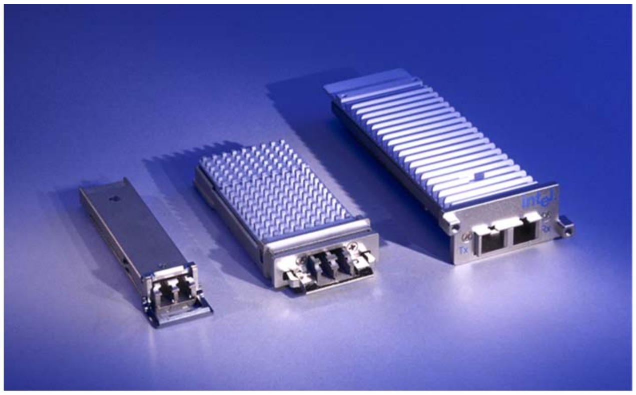

Figure 3: Next-generation 10 Gb/s enterprise optical transceivers: (from left) XFP, XPAK/X2, XENPAK. These modules are electrically hot-pluggable and optically pluggable.

In general, not all the switches' ports are populated with transceivers when they are shipped to customers. The customer has the option to buy these transceiver modules as the demand for ports increases. It also gives the customer the choice of optics: MMF or SMF. Therefore, these modules have been designed to be pluggable.

The choice between the different 10 Gb/s form-factor optical transceiver packages is guided by reach, cost, and thermal and size constraints and requirements.

Table 1: Summary of different form-factor 10 Gb/s optical transceiver packages

MSA	XENPAK	XPAK/X2	XFP
MSA Date	March 2002	March 2003	April 2003
Application	Enterprise switch	Enterprise switch NIC Storage	Telecom Datacom
Electrical Interface	4 bit XAUI	4 bit XAUI	1 bit XFI
Optical interface	SC pluggable	SC or LC pluggable	LC pluggable
Dimension	4.8x1.4x0.7	2.7x1.4x0.4	
Max Power	11W	5W	3W

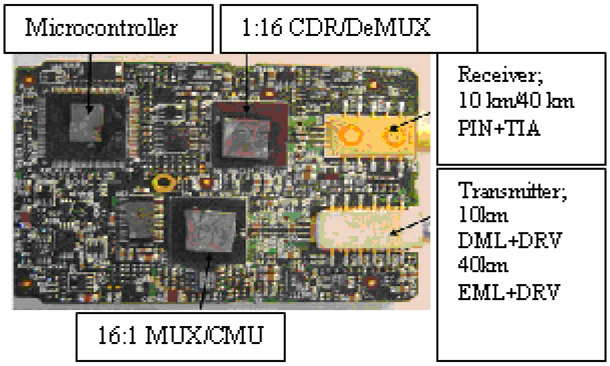

Figure 4: Intel® TXN13220 FR-4 printed circuit board showing optical modules, Mux/DeMux, and microprocessor

OPTICAL TECHNOLOGY TRENDS IN THE ENTERPRISE

Enterprise network topology can be divided into four main categories: horizontal cabling, vertical cabling, data center interconnects, and campus backbone. Table 2 shows the use of various technologies in each of these topologies based on distance and speed. Within a data center, the LAN, which is based on the Ethernet protocol, is based mostly on copper interconnects operating at speeds of up to 1 Gb/s for all distances below 100 m. Some MMFs at 850 nm are also used at speeds of up to 10 Gb/s. However, within a SAN, which is based on the Fiber Channel (FC) protocol, only MMFs at 850 nm are used for transmission rates ranging form 1 Gb/s to 10 Gb/s. Optical interconnects based on SMFs at 1310 nm are typically used within a campus-size network with transmission distances up to 10 km and rates of up to 10 Gb/s. All the optical interconnects used in enterprise networks today are box-to-box connections typically between a server to a switch connection and a switch to a switch connection. The most popular optical modules are SSF and SFP form-factor modules operating at speeds of up to 4 Gb/s, while the 1 Gb/s GBIC for LAN is still shipping in volume. For the 10 Gb/s modules, the XENPAK is the most commonly used today within enterprise networks. However, the new optical module designs are trending towards smaller form-factor modules such as XPAK and XFP.

Table 2: Used technology in each of the enterprise network topologies

Segment	Distance	Speed	Technology
Horizontal	100 m	10/100/1000 Mb/s	Copper
Vertical	300 m	10/100 Mb/s 1/10 Gb/s	Copper MMF
Data center	100 m	10/100Mb/s 1Gb/s-10Gb/s LAN 1/2/4 Gb/s SAN	Copper Copper/MMF MMF
Campus	2-10 km	1- 10 Gb/s	SMF

In today's enterprise networks, copper-based interconnects dominate the box-to-box connections due to their cost advantage over comparable optical interconnects. For example, in the 1 Gb/s LAN (i.e., Ethernet-based) market, Dell'Oro in January 2004 was forecasting the shipment volume to grow from about a total of 13 million switch ports in 2003 to 158 million ports in 2008 [11]. Out of the total number of switch ports, the number of 1 Gb/s copper-based interconnects is expected to grow from about 7 million in 2003 to about 141 million in 2008. This anticipated growth of copper-based interconnects is driven by the deployment of 1 Gb/s to the desktop. The average selling price of copper-based switch ports is $181 in 2003. For the lowest cost segment of the copper-based 1 Gb/s switch port market, the average selling price is forecasted to decline from about $59 to $9 over the same period. In comparison, out of the total number of switch ports, the 1 Gb/s optical interconnects market is expected to grow from about 6 million ports in 2003 to 17 million ports in 2008 with the average selling price about $416 in 2003 [11]. The market for the SAN, which is based on the FC protocol standard, is planned to transition from optical interconnects operating at 2 Gb/s to 4 Gb/s during the 2005-2006 period. The forecasted volume for the FC-based switch ports is expected to grow from 1.8 million of 2 Gb/s ports in 2003 to 6.4 million of 4 Gb/s ports in 2007, according to the IDC forecast published in August 2003 [12]. Over the past year, the average selling price of optical 10 GbE switch ports has been dropping from about $30K to about $10K per port today. The average selling price of these 10 GbE switch ports must be below $3K per port for a wide deployment within enterprise networks.

NEXT-GENERATION OPTICAL INTERCONNECTS

In general, interconnects can be divided into four main categories: box-to-box, board-to-board, chip-to-chip, and on-chip interconnects. Today, optical interconnections in the enterprise are mostly used for the box-to-box interconnects. The optical transceivers are either plugged into high-capacity Ethernet or FC-based switches, or placed on a NIC/HBA, which are implemented into servers. Will optical interconnects move into the other categories? A brief overview of each optical interconnect category is provided in the next subsections.

Board-to-Board Interconnects

The move of the communications industry, both telecom and datacom toward a Modular Communication Platform (MCP) favors the deployment of bladed architecture. Functions that were traditionally housed in a standalone box, such as between servers and switches, have started to be implemented into a board-level form-factor, which is called a blade, and they plug into a common chassis. Therefore, interconnect opportunities at the board level are becoming more important. As the transmission speed increases, copper-based interconnects are facing technical challenges in terms of speed, reach, EMI, and routing.

Backplanes are potential applications for optical interconnects. These are point-to-point or point-to-multipoint high-speed interconnects with typical lengths of under 1m. The key advantages of optical backplane interconnects are low-crosstalk among the optical signals, and their large bandwidth. However, most of today's optical backplanes are more like patch panels rather than replacements for backplanes. Many different optical technologies have been demonstrated including polymer waveguides integrated on Si, planar light wave circuit interconnects, and fiber ribbon arrays integrated with VCSELs and photodiodes. However, none of these optical technologies today have displaced copper interconnects outside of some niche applications. The transition to optical backplanes might be induced by the accumulating technical challenges of electrical interconnects. However to be widely adopted, optical interconnects must be able to advance toward a smaller form-factor with a lower power consumption at a lower cost. Meeting these requirements is critical to the technology evolution of optical interconnects from box-to-box to board-to-board to chip-to-chip. Several new technical breakthroughs will need to occur in order to meet these challenging requirements.

Chip-to-Chip Interconnects

Extrapolation of Moore's law shows that microprocessors are expected to be clocked at about 10 GHz by the end of the decade [13, 14]. Consequently, it is becoming extremely difficult to route enough bandwidth through a PCB or a module using the existing electrical wires. It has been shown that frequency-dependent loss for copper traces on FR4 circuit boards rapidly rises above 1 GHz, reducing the Signal-to-Noise Ratio (SNR) and introducing timing errors. In addition, the resultant high-frequency crosstalk among the different copper traces limits the wiring density on the circuit board. High-speed short-distance ($L < 10$ cm) chip-to-chip optical interconnects have several advantages over copper interconnects. Optical interconnects are low-loss interconnects with a large transmission bandwidth. Another key advantage is their inherent immunity to Electro-Magnetic Interference (EMI). The density of copper traces on FR4 boards is constrained by these EMI and electrical crosstalk problems. Over the last 20 years, many different optical technologies were demonstrated to overcome the electrical bottleneck [14, 15]. However, their relative high costs due to implementation complexity and use of exotic materials made them unsuitable for high-volume manufacturing and thus prevented the adoption of these technologies.

On-Chip Interconnects

The design of on-chip electrical interconnects is becoming increasingly difficult given the continuous growth in the complexity of integrated circuits operating at multi-GHz. Can on-chip optical interconnects potentially solve these issues? On-chip point-to-point and point-to-multipoint optical interconnects with typical lengths under 1 cm are potentially attractive because of the following reasons: (a) they decrease the existing electrical interconnect delays, (b) provide a higher bandwidth to keep pace with the speed of transistors, (c) reduce electrical power consumption, and (d) minimize sensitivity to EMI. The potential primary application of on-chip optical interconnects in microprocessors, for example, are in high-speed signaling and clock distribution. For on-chip signal distribution, four key benchmark parameters are typically used: signal delay normalized by clock cycle, available bandwidth per unit area or bandwidth density, bandwidth density/delay ratio, and cost. For on-chip clock distribution, the critical parameters are timing, skew, and jitter. The primary challenges for implementing these optical interconnects is the integration of multiple VCSEL and photodiode arrays with their corresponding drivers and TIAs, and on-chip light coupling into optical waveguide arrays over the entire chip. Currently, such optoelectronic integration is not only in its early stages of development, but it is also very expensive relative to copper-based interconnects with limited or no performance advantages. In order to take advantage of on-chip optical interconnects, today's microprocessor architecture might need to evolve from a single superscalar chip to a mesh of optically interconnected processors with their associated memories. The WDM scheme could be used, for example, to send multiple wavelengths in the same optical waveguide to significantly increase the overall communication bandwidth.

Silicon-Based Optical Interconnects

In order to enable the chip-to-chip or on-chip optical interconnects, silicon-based optical components should be developed. As previously noted, optical components such as tunable WDM filters, photodiodes, optical waveguides, and electronic components such as laser drivers and TIA circuits are all based on different materials and technologies precluding them from an optoelectronic monolithic integration. The attractiveness of silicon-based optical interconnects is the potential integration with CMOS integrated circuits for high-volume manufacturing. The first prototype of such silicon-based optical components includes thermally tunable WDM Bragg filters, high-speed optical MZI modulators and a Si/Ge high-speed photodiode [16].

A narrow-band Bragg grating filter has been made in silicon waveguides with alternating polycrystalline/crystalline layers. The Bragg grating reflects only the wavelengths that satisfies the Bragg condition [1–4]

$$\lambda_B = 2n \cdot \Lambda / m$$

where n = 3.46 in the effective refractive index of the silicon waveguide, Λ = 2.445 μm is the Bragg grating pitch, and m = 11 is the Bragg grating order. Using the strong thermo-optics effect in silicon, these Bragg grating filters can be made tunable [17]. Such WDM Bragg grating filters with a 3-dB bandwidth of 100 GHz and 200 GHz, insertion loss of about 4 dB and tunability over 12 nm for 100 °C temperature variation were demonstrated. These tunable Bragg filters can be used as channel filters in a low-cost WDM communication system.

Silicon-on-Insulator (SOI) optical modulators based on current injection had a modulation bandwidth of only 20 MHz. Consequently, these devices were not suitable for today's high-speed communication networks.

Recently, an MZI silicon optical modulator with a modulation bandwidth of 2.5 GHz around 1550 nm was demonstrated [18]. The high-speed operation was achieved by using a novel phase shifter design based on a Metal-Oxide-Semiconductor (MOS) capacitor embedded in a passive silicon waveguide with doped and undoped polysilicon regions in MZI configuration. The accumulated charges in the MOS capacitor induce fast refractive index changes in the silicon waveguide due to the free carrier plasma dispersion effect [19]. The 1.5 cm long silicon MZI modulator had an extinction ratio of more than 16 dB with an applied peak-to-peak voltage of about 7.7 volts.

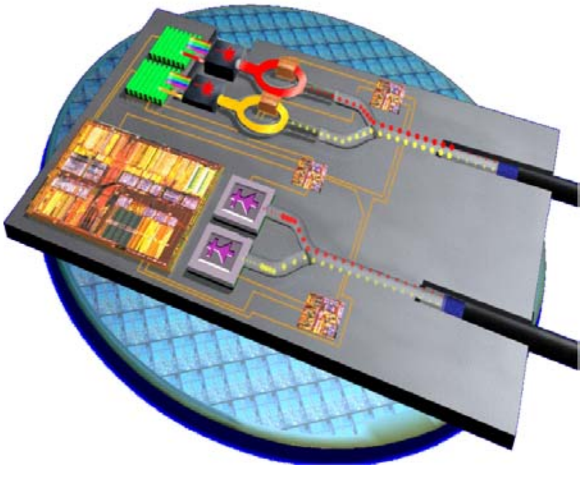

Figure 5: Future vision of hybrid optical integration of a four-channel WDM optical transceiver with silicon photonic components and conventional CMOS drivers

However, the demonstrated modulator had a total of 12 dB insertion loss, which can potentially be reduced to about 4–5 dB if the polysilicon is replaced with a single-crystal silicon and tapered waveguides are used to efficiently couple the light in and out of the modulator through a diabatic transformation of the optical mode. The development of tapered waveguide technology for this high-index contrast optical system is one of the key challenges to developing a hybrid integration of these devices.

Such SOI-based modulators and Bragg filters can lead to an integrated optical transceiver module on a single die, combining multiple wavelength sources, modulators, and WDM multiplexer/de-multiplexers with the corresponding drive electronics. This concept is illustrated through an artist vision in Figure 5 which shows a hybrid integration of a four-channel WDM optical transceiver with silicon photonic components such as modulators, filters, and multiplexer/de-multiplexers and conventional CMOS-based drivers on the same silicon die.

TECHNICAL CHALLENGES AND CONCLUSION

The costs of high-speed optical interconnects has been reduced by more than an order of magnitude while their performance has been significantly improved over the last five years. Tremendous progress has been made in the development of cost-effective packaging technologies for optical modules operating at transmission rates up to 10 Gb/s. Four different MSA form-factor packages for 10 Gb/s optical transceivers in enterprise networks, namely, XENPAK, XPAK, X2, and XFP have been created, trending toward smaller form factors. In the meantime, the SFF and SFP, which are the smallest form-factor modules on the market, are trending to higher speed capabilities from 1/2 Gb/s today toward 4 Gb/s. In today's enterprise networks, the optical transceivers are used for box-to-box connections. The optical transceivers are either plugged into high-capacity Ethernet or FC-based switches, or placed on a NIC/HBA, which are plugged into servers. In terms of volume shipped, copper-based interconnects dominate the box-to-box connections due to their cost advantage over comparable optical interconnects until at least 2008. However, the number of either Ethernet-based or FC-based server-to-switch and switch-to-switch optical interconnects is expected to grow in the coming years.

To be widely adopted, optical interconnects must be able to advance toward a smaller form-factor with a lower power consumption at a lower cost. Meeting these requirements is critical to the technology evolution of optical interconnects from box-to-box to board-to-board, to chip-to-chip. Several new technology breakthroughs will need to occur in order to meet these challenging requirements.

ACKNOWLEDGMENTS

The authors would like to thank Dr. Mario Paniccia for technical discussions and support. Pierre would like to acknowledge Pete Kirkpatrick, Marc Epitaux, Eric Zbinden, Mark Finot, Marc Verdiell, and Jeff Bennett for their technical support.

REFERENCES

[1] Agrawal, G. P., and Dutta, N. K., *Long-Wavelength Semiconductor Lasers*, Van Nostrand Reinhold, New York, 1986.

[2] Agrawal, G. P., *Fiber-Optic Communication Systems*, John Wiley & Sons, New York, 1997.

[3] Goff, R., *Fiber Optics Reference Guide,* Focal Press, Third Edition, 2002.

[4] Ovadia, S., *Broadband Cable TV Access Networks: From Technologies to Applications*, Prentice Hall PTR, Upper Saddle River, NJ, 2001.

[5] Suzuki, M., and Otani, T., "High-Speed (40–160 Gb/s) WDM Transmission in Terrestrial Networks," 2003 *OFC Technical Digest*, pp. 741–742, paper FN1.

[6] Corning SMF-28™ Optical Fiber Product Information (2002).

[7] Wooten, E. L., "A Review of Lithium Niobate Modulators for Fiber-Optic Communications Systems," *IEEE Journal on Selected Topics in Quantum Electronics* **6**, pp. 69–82, 2000.

[8] Ido, T. et al., "Ultra-High-Speed Multiple-Quantum-Well Electro-Absorption Optical Modulators with Integrated Waveguides," *IEEE Journal of Lightwave Technology* 14, pp. 2026–2034, 1996.

[9] Takemoto, A., Watanabe, H., Nakajima, Y., Sakakibara, Y., Kakimoto, S., Yamashita, J., Hatta, T., and Miyake, Y., "Distributed Feedback Laser Diode and Module for CATV Systems," *IEEE Journal of Selected Areas in Communications* 8, 1359–1364, 1990.

[10] Longford, A., and Radloff, B., "The Use of Pre-Molded Leadframe Cavity Package Technologies in Photonic and RF Applications," 27th Annual IEEE/SEMICON conference proceedings, 348–352, July 17–18 (2002).

[11] Dell'Oro Group Ethernet Switch 5-Year Forecast, January 2004.

[12] IDC Worldwide Fibre Channel Hub and Switch Forecast and Analysis, 2003-2007 Aug. 2003, Doc #29882 Study

[13] Huang, D., Sze T., Landin A., Lytel R., Davidson, H. L., "Optical Interconnects: Out of the Box Forever?," *IEEE J. of Selected Topics in Quantum Electronics* 9, 614–623, 2003.

[14] Collet, J. H., Caignet, F., and Litaize, D., "Performance Constraints for Onchip Optical Interconnects," *IEEE Journal of Selected Topics in Quantum Electronics* 9, 425–432, 2003.

[15] Goodman, J. W., Leonberger, F. J., Kung, S.-Y., and Athale, R. A., "Optical Interconnects for VLSI Systems," *Proceedings of IEEE* 72, 850–866, 1984.

[16] Colace, L., Masini, G., and Assanto, G., "Ge-on-Si Approaches to the Detection of near-Infrared Light," *IEEE Journal of Quantum Electronics* 35, 1843–1852 (1999).

[17] Soref, R. A., "Silicon-based Optoelectronics," *Proceedings of IEEE* 81, 1687–1706, 1993.

[18] Liu, A., Jones, R., Liao L., Samara-Rubio D., Rubin, D., Choen O., Nicolaescu R., and Paniccia M., "High-Speed Silicon Optical Modulator Based on a Metal-Oxide-Semiconductor Capacitor," *Nature* 427, 615–618, 2004.

[19] Soref, R. A., and Bennett, B. R., "Electro optical effects in Silicon," *IEEE Journal of Quantum Electronics* 23, 123-129, 1987.

AUTHORS' BIOGRAPHIES

Pierre Herve is currently director of Enterprise Strategic Planning at Intel in the Communication Infrastructure Group. He has been with Intel since 1996 and has held various marketing positions including director of optical strategy, wireless senior marketing manager, and CATV product manager. Before joining Intel, Pierre managed an online services provider based in Northern California for five years. He graduated with a M.S. degree in Telecommunications from the Institut National des Telecommunications in 1988. His email is pierre.m.herve at intel.com.

Shlomo Ovadia received a Ph.D. degree in Optical Sciences from the Optical Sciences Center, University of Arizona in 1984. After two years as a postdoctoral fellow at the University of Maryland, Shlomo held various technical positions at IBM, Bellcore, and General Instruments. In 2000, Shlomo joined Intel in San Jose, California, as principal architect in CTG, where he was leading the effort on the architecture, design, and development of optical burst switching in enterprise networks. Currently at Intel Communication Group, Shlomo is responsible for WiMAX interoperability and certification for IEEE 802.16d/e wireless modems. He is the author of a recently published book titled *Broadband Cable TV Access Networks: From Technologies to Applications* (Prentice Hall, 2001). He is a Senior Member of IEEE/LEOS/COMSOC with more than 70 technical publications and conference presentations. He is the holder of 35 patents, and his personal biography is included in the Millennium edition of *Who's Who* in Science and Engineering (2000/2001). His email is shlomo.ovadia at intel.com.

For further information visit:

developer.intel.com/technology/itj/index.htm