



Intel[®] Technology Journal

Optical Technologies and Applications

10 Gb/s Optical Transceivers: Fundamentals and Emerging Technologies

10 Gb/s Optical Transceivers: Fundamentals and Emerging Technologies

Peter Kirkpatrick, Intel Communications Group, Intel Corporation
Wei-chiao Fang, Intel Communications Group, Intel Corporation
Henrik Johansen, Intel Communications Group, Intel Corporation
Benny Christensen, Intel Communications Group, Intel Corporation
Jesper Hanberg, Intel Communications Group, Intel Corporation
Martin Lobel, Intel Communications Group, Intel Corporation
Tom Mader, Intel Communications Group, Intel Corporation
Song Shang, Intel Communications Group, Intel Corporation
Craig Schulz, Intel Communications Group, Intel Corporation
Doug Sprock, Intel Communications Group, Intel Corporation
Marc Verdiell, Intel Communications Group, Intel Corporation

Index words: fiber optics, optical transceiver, 10 Gb/s, Dense Wavelength Division Multiplexing (DWDM), tunable laser, Electronic Dispersion Compensation (EDC)

ABSTRACT

Continued demand for higher bandwidth in data networks along with an industry need for standardized network components has driven the need for 10 Gb/s optical transceivers into the long haul, metro core, and enterprise/storage environments. In response to this demand, the functionality and performance of these modules have increased significantly, while the size, cost, and power consumption have been dramatically reduced. While each transceiver application has unique system-level requirements, performance criteria, and cost sensitivities, the structure of each transceiver is much the same. This paper presents a discussion of the basic transceiver architecture including analog and digital electronic ICs and process choices, optical components and packaging, and control, monitoring, and interface circuitry. With each component, the critical performance criteria and cost-reduction opportunities are highlighted. More detail is presented on a full C-band wavelength tunable long-haul optical transceiver for Dense Wavelength Division Multiplexed (DWDM) systems, including the DWDM system requirements and design issues that arise in amplified, high-power, multi-channel links. The functionality and performance of a tunable 10 Gb/s transceiver are explored in more depth with a focus on transmitter and receiver front-end design. Finally, an application of Electronic Dispersion Compensation

(EDC) is presented, and we show how it can be used to compensate for fiber limitations to enable extended 10 Gb/s links over modal bandwidth limited legacy multimode fiber in enterprise systems.




Figure 1: 10 Gb/s optical transceivers for DWDM line side and metro client-side interfaces. Pictured from the left are the 300-pin, XFP, and small form factor 300-pin MSA modules.

INTRODUCTION

Optical transceivers operating at line rates of 10 Gb/s have matured rapidly over the last few years and are currently available in a wide variety of form factors, each addressing a range of link parameters and protocols. These form factors are the result of Multi-Source

Agreements (MSAs) that define common mechanical dimensions and electrical interfaces. The first MSA was the 300-pin MSA in 2000, followed by XENPAK, X2/XPAK, and XFP. Each of the transceivers defined by the MSAs have unique advantages that fit the needs of various systems, supporting different protocols, fiber reaches, and power dissipation levels. Figures 1 and 2 show 10 Gb/s transceivers for client-side telecom and enterprise interfaces, respectively. As the underlying technologies have matured, the optical and electrical performance of these modules has been significantly improved, while the power dissipation and cost have been reduced dramatically.



Figure 2: 10 Gb/s optical transceivers for enterprise interfaces. Pictured from the left are the XFP, X2/XPAK, and XENPAK MSAs.

In the architecture section of this paper, the basic architecture of all 10 Gb/s transceivers is presented, highlighting the functionality and critical performance parameters of each functional block. Design considerations such as IC process and optical packaging technologies are also discussed.

In the DWDM section, details on the design challenges and performance of a 10 Gb/s full C-band wavelength tunable optical transceiver are presented. Issues that arise in amplified, high-power, multi-channel links are discussed in terms of design choices and performance criteria.

Finally, the fundamentals of Electronic Dispersion Compensation (EDC) are presented, which uses adaptive electrical filtering techniques to compensate for limitations incurred during fiber propagation. The application of EDC to the problem of extending 10 Gb/s links using 1310 nm directly modulated lasers over low-bandwidth Multi-Mode Fiber (MMF) is explored in more detail.

10 GB/S TRANSCEIVER ARCHITECTURE

The basic architecture and interconnection between components in a 10 Gb/s optical transceiver are shown in Figure 3. This basic architecture holds for all of the 10 Gb/s transceiver Multi-Source Agreements (MSAs), independent of form factor, electrical interface, or fiber reach. Figure 4 shows the actual printed circuit board from a small form factor 300-pin MSA transceiver.




Figure 3: 10 Gb/s optical transceiver architecture showing functional components and interconnection

Electrical Interface

The electrical interface provides input/output data transfer to the host card, various clocking channels, control and monitoring channels, as well as DC power and ground connections. This interface can take the form of a socket that plugs perpendicularly into the plane of the host board, as in the 300-pin MSA [1], or it can take the form of a board-edge connection that mates to a socket in the plane of the host board. The latter is the case for XENPAK [2], X2/XPAK [3,4], and XFP [5], which provide front-panel pluggability at the host system level. Along with this benefit comes the added requirement of hot-pluggability and inrush current management.




Figure 4: Small form factor 300-pin optical transceiver PCB showing optical modules with integrated driver/TIA, PMA (1:16 Mux/DeMux), and microcontroller

The width of the data bus varies between the different MSAs. The first to emerge was the XSBI (based on SFI-4) interface of the 300-pin MSA, with 16-bit differential buses for input and output data [6]. Each I/O channel transmits or receives data at 1/16 the line rate (e.g., 622.080 Mb/s with a SONET line rate of 9.953 Gb/s). The XAUI interface on XENPAK and X2/XPAK provides a 4-bit differential bus. In this interface, each channel operates at 3.125 Gb/s, for a total bandwidth of 12.5 Gb/s. Multiple encoding/decoding steps on the XAUI signals result in the Ethernet line rate of 10.3125 Gb/s, which we discuss further in the Physical Medium Attachment (PMA) section. Finally, for XFP, the XFI interface provides a 1-bit differential signal at the line rate, providing a truly protocol-independent interface for 10 Gb/s communication.

Clocking schemes vary between the different interfaces. The clocking used in SFI-4 systems can vary, but generally a synchronous clock (1/16 or 1/64 of the line rate) is provided with the input data, and a reference clock is provided to keep the receiver Clock and Data Recovery (CDR) within locking range of the input data rate. The clock recovered by the receiver CDR is provided along with the output data. In XAUI and XFP transceivers, the host board simply provides a reference clock to keep the CDRs in the module within locking range of the data rate.

The electrical interface also provides DC connections to the host board DC power supplies. Care must be taken to minimize power dissipation given the different voltage and current levels available in various system designs.

Control System

The control systems in modern 10 Gb/s transceivers are generally implemented using a microcontroller. This device performs many functions that previously had been implemented using analog hardware. The controller sets control parameters for the Physical Medium Attachment (PMA), transmitter, and receiver, which may vary over time and temperature or when the host system changes the link configuration (i.e., loopback modes). The controller also provides a two-wire interface such as I²C so the host board can set control parameters and read the status registers where monitor values are stored.

Physical Medium Attachment

The core electrical functionality of a 10 Gb/s transceiver is contained in the PMA device. Transceivers at lower data rates generally do not include this functionality, because maintaining robust signal integrity is not as difficult at lower speeds. PMA devices take the form of clock multiplier/multiplexer (MUX/CMU) and clock and data recovery/demultiplexer (CDR/DeMUX) in the 300-pin modules. The MUX/CMU interleaves the 16-channel

data bus into a serialized data stream at the line rate, clocked by a multiplied version of the input clock. This data stream is used to modulate the optical transmitter. The CDR/DeMUX provides the complementary functionality on the receive side. For the two types of PMA, there are a wide range of features and performance parameters that are critical to system-level quality. Here, we focus on the basic functionality and jitter characteristics. The PMA output waveform quality has a large impact on optical transmit eye quality, and PMA sensitivity has a large impact on optical receive sensitivity. The system-level impacts of these characteristics are covered in the optical device section below.

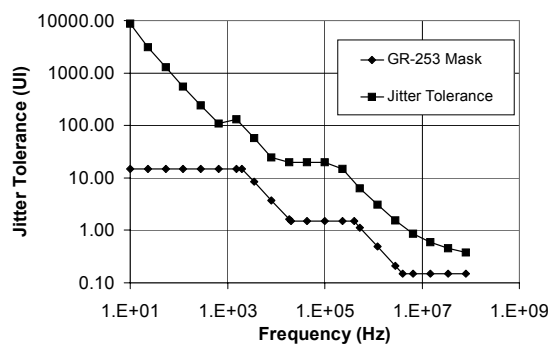


Figure 5: Jitter tolerance performance of a typical 300-pin small form factor 40 km transceiver. Note the stepped mask indicating the specification limit.

In SONET systems [7], jitter performance is critical to maintaining error-free performance over links involving several network elements [8]. There are three jitter parameters of interest in these systems: transmitter jitter generation, receiver jitter tolerance, and receive-to-transmit jitter transfer. For all SONET network elements, the jitter generation must have significant margin on the specification limit of 100 mUIpp (UI denotes a unit interval, or one bit period), ideally 50 mUIpp. For a receiver, robust jitter tolerance is critical to performance of the network. This is the amount of sinusoidal jitter that causes a 1 dB penalty in the sensitivity of the receiver. Figure 5 shows the worst-case jitter tolerance performance of the 300-pin small form factor transceiver along with the specification mask. When a transceiver is used as a simple repeater, the jitter transfer characteristic is also important to network reliability. This is the amount of jitter transmitted at a specific frequency given a specified level of jitter into the transceiver. The worst-case jitter transfer performance of a 300-pin small form factor transceiver is shown in Figure 6. The peaking is less than 0.1 dB at low frequencies, and at higher frequencies the transfer rolls off to stay below the specification limit.

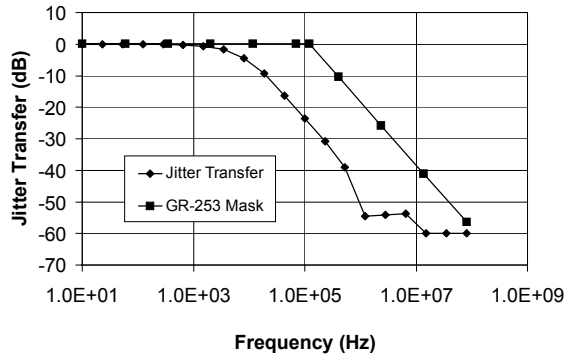


Figure 6: Jitter transfer performance of a typical 300-pin small form factor 40 km transceiver. The specification limit is indicated on the plot.

Although Ethernet standards do not specifically define jitter parameters, the jitter generation and tolerance of these devices are still critical in Ethernet systems [9]. Jitter generation has an impact on the Transmitter and Dispersion Penalty (TDP), where the transmission performance over a fiber link of a device under test is compared with a reference transmitter. Any degradations due to waveform imperfections or jitter generation add to this penalty. Jitter tolerance has an impact in the Stressed Receive Sensitivity (SRS) test, where receiver sensitivity is measured with an eye that contains stress from sinusoidal jitter, as well as filtering and vertical sinusoidal closure. The TDP and SRS performance of a 40 km XENPAK transceiver is shown in Table 1. As 300-pin modules are often used in Ethernet systems, the TDP and SRS performance of the 40 km 300-pin small form factor module is shown for comparison.

Table 1: 10 Gb/s Ethernet 40 km (10GBASE-L) optical transceiver performance

SPECIFICATION	300-PIN 40 KM		XENPAK 40 KM		
	Performance	Margin	Performance	Margin	
SRS	<-11.3 dBm	-15.1 dBm	3.8 dB	-15.4 dBm	4.2 dB
TDP	< 3 dB	0.6 dB	2.4 dB	1.1 dB	1.9 dB

Data is worst-case performance over temperature (-5..70°C) and power supply variation (±10%).

The XAUI PMA used in Ethernet and Fiber Channel systems also has MUX/CMU and CDR/DeMUX functionality for the 4-lane interface. The definition of the XAUI interface provides the additional benefits of individual CDRs on each lane as well as several bits of deskew functionality to provide robustness to layout and routing issues. The architecture also defines encoding and decoding overhead to the interface, to provide error detection, link status, and DC balance. On the 4-lane side,

each lane operates at a line rate of 3.125 Gb/s, for an aggregate bit rate of 12.5 Gb/s. The 8B/10B coding on this interface provides DC balance at a 25% overhead, for an actual data rate of 10 Gb/s. The signal is then scrambled, and 64B/66B coding is used to add link status, line/block coding, and scrambler synchronization on the optical interface, bringing the serial line rate to 10.3125 Gb/s.

For XFP transceivers, the PMA consists of simple CDRs for both the transmit and receive paths. As these modules work in both SONET and Ethernet systems, all of the performance parameters above apply. They must provide robust SONET jitter performance as well as TDP and SRS. This is particularly difficult in SONET systems as the output of an XFP transmitter must meet the 100 mUIpp jitter generation spec, while the chip providing the data signal to the XFP unit is allowed to have 50 mUIpp of jitter.

Optical Devices and Physical Medium Dependent ICs
Depending on the reach requirements of a fiber link, different optical transmit and receive devices are required. Although there are alternative devices and technologies emerging, devices that are technically established and available in the market are discussed here.




Figure 7: 850 nm multimode VCSEL optical eye diagram at 10 Gb/s from an XFP module

For Very Short Reach (VSR) applications in the enterprise space, such as within-building or within-campus LAN and SAN connections where multimode fiber (MMF; 50 or 62.5 μm core diameter) is present, 850 nm GaAs/AlGaAs Vertical Cavity Surface Emitting Laser (VCSEL) transmitters and GaAs pin photodetector-based receivers are sufficient for 10 Gb/s on links of 25 to 300 m, depending on fiber quality. The low drive current and high efficiency of a VCSEL enable a transceiver with very low power dissipation. As the drive current and voltages are low, generally less than 20 mApp, a low-power CMOS driver can be used for direct modulation of

the VCSEL current. The bias current and driver swing may be controlled over temperature to compensate for changes in the characteristics of the device in uncooled operations. As these devices are only used in Ethernet links, they are generally operated at lower extinction ratios (> 3.5 dB). A typical 850 nm VCSEL transmit eye diagram is shown in Figure 7.




Figure 8: 1310 nm directly modulated DFB optical eye diagram at 10 Gb/s from small form factor 300-pin 10 km module

For Longer Reach (LR) links (7-20 km), a single-mode fiber (SMF; 9 μm core diameter) is required, along with an InP-based single-mode Distributed Feedback Laser (DFB) operating at 1310 nm. This wavelength is chosen to match an operating region in the fiber that has close to zero chromatic dispersion (an impairment which otherwise causes pulses to expand during propagation). To satisfy SONET extinction ratio specifications (> 6 dB), a DFB has higher current swing requirements than a VCSEL, which drives the need for higher current GaAs or SiGe drivers, lower impedance (25 Ω) interfaces, and alternative modulation schemes to save power, improve yields, and lower cost. Care must be taken to compensate for temperature variations in uncooled formats. A typical 1310 nm, 10 km DFB eye diagram is shown in Figure 8.

For metro area network access, Extended Reach (ER) links (40-80 km) may be used over SMF. These links are traditionally addressed by using externally modulated lasers, or EMLs, at 1550 nm. These devices, which monolithically combine a DFB laser with an Electro-Absorption Modulator (EAM) on an InP substrate, can be tuned to operate with low chirp (change of wavelength under modulation) to enable extended propagation in the low attenuation but higher dispersion 1550 nm window. These devices are generally modulated with GaAs or SiGe drivers to provide ample voltage swing and offset to achieve high extinction ratios (> 9 dB) and high-quality eye diagrams. EMLs are operated in controlled temperature environments to ensure performance, the cost

of which is increased power dissipation from the Thermo-Electric Cooling (TEC) element and control circuit. A typical 1550 nm, 40 km EML eye diagram is shown in Figure 9.




Figure 9: Small form factor 300-pin 40 km module 1550 nm optical eye diagram at 10 Gb/s

This type of eye quality must be achieved without sacrificing dispersion penalty performance, or difference between back-to-back sensitivity and post-dispersion sensitivity. High-quality 40-80 km transmit modules have low-dispersion penalty (< 1 dB).

Optical Receivers

Receivers for 10 Gb/s links are based on two types of optoelectronic devices, InP or GaAs PIN (Positive-Intrinsic-Negative doped structure) photodiodes or InP Avalanche Photodiodes (APD) that convert light into electrical current. The fundamental performance criteria for a receiver is its optical sensitivity, which is the optical power required to maintain a link Bit Error Rate (BER) of 10^{-12} . A typical sensitivity curve is shown in Figure 10.

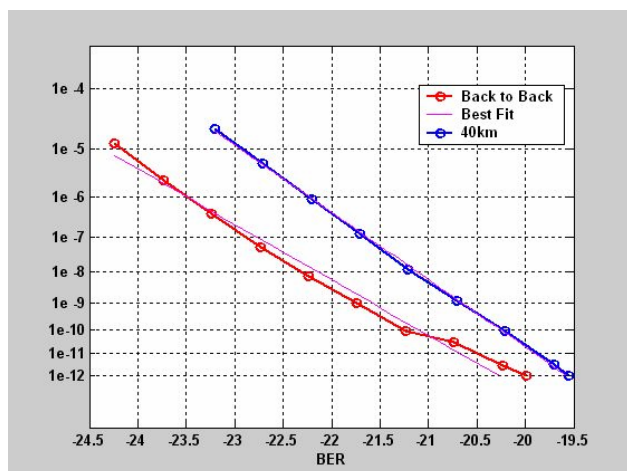


Figure 10: Sensitivity curves from a typical 300-pin small form factor transceiver. The BER is plotted in error function space to yield a straight line.

At these data rates, receivers integrate Transimpedance Amplifiers (TIA) to convert the photocurrent into a voltage swing that can be provided to a decision circuit such as a CDR. The shorter reach links above, from 300 m at 850 nm up to 40 km at 1550 nm, all use pin-based receivers. Unamplified links at 80 km use the additional gain of an APD to meet more stringent sensitivity requirements. Generally, a Limiting Amplifier (LIA) is integrated into the TIA to provide high gain and high sensitivity. There are cases, discussed later, where high gain is traded for linearity. Depending on the noise and gain requirements set by the sensitivity level of a particular link, TIAs can be designed in GaAs, SiGe, or CMOS processes.




Figure 11: Optical module packaging options. Upper left shows a standard cooled “butterfly” can with coaxial RF interface. Upper and lower right show cooled and uncooled modules, respectively, with coplanar RF interface. Lower left shows uncooled TO can with coplanar RF interface.

Optics Packaging

Depending on the performance, cooling, and board interface requirements of a particular transceiver design, several different optical packaging technologies are employed, as shown in Figure 11. Traditionally, due to the difficult problem of achieving stable high optical coupling in a reliable manner, “butterfly” style packages have been used. Although they are large and costly, these packages provide a stable, hermetic environment for transmitters, TECs, and receivers, along with a coaxial RF interface to ensure high-quality signal integrity. Smaller, lower cost cooled packages with coplanar RF interfaces have been developed to enable smaller form factor transceivers without sacrificing performance. Uncooled designs are also smaller and cost a lot less; although, their use is limited to devices that can operate over a wide temperature range such as receiver circuitry and uncooled transmitters. Finally, coaxial TO-can-style optics with either coplanar ceramic or glass feedthrough RF interfaces are emerging to provide a very low-cost package that still provides adequate signal integrity for 10 Gb/s operation. While miniaturizing and reducing the cost of these optical modules, it is important to maintain performance and the standards of reliability that have been met by modules qualified to the Telcordia requirements in GR-468 [10]. It is critical that new optical packaging technologies increase yields and the quality of performance and reliability while decreasing the cost of the components, simplifying their manufacture, and reducing testing.

With the basic architecture of the 10 Gb/s optical transceivers presented here, optical communication system vendors have a wide array of transceivers available to meet the specific needs of individual system designs. Depending on the requirements of a particular system design, various form factors, fiber reaches, and performance levels are available in standardized modules.

Two emerging technologies that are influencing 10 Gb/s transceiver design are discussed in the following sections. The first, a wavelength tunable transceiver, enables next-generation DWDM systems. The second, EDC, can compensate for signal degradations due to fiber limitations.

DWDM DESIGN CHALLENGES

Introduction

The recent downturn in the telecom market has forced telecommunications carriers to put enormous pressure on equipment providers to reduce system costs in the core backbone, where DWDM transport systems are typically employed. As a result, equipment providers are beginning to outsource their line-side optics to transceiver vendors. Cost pressures will continue to push standardized

transceiver-based solutions into next-generation products as the DWDM market continues to grow. However, standardizing DWDM line-side optics is very difficult given the very stringent and proprietary requirements specific to each system provider. Here, we present the technologies employed in a 10 Gb/s DWDM transceiver to guarantee robust transmission.

DWDM transport systems are made possible by one key invention, the Erbium-Doped Fiber Amplifier (EDFA). An EDFA boosts the optical power without the need for an electrical regeneration, and thus it manages the link losses between network nodes. It simultaneously amplifies any optical channel within its gain-bandwidth, and this is why optical transport system designers densely pack channels together. Telecommunications carriers can install EDFAs at repeater sites between each fiber span, usually every 20-80 km. The benefits of EDFAs include the ability to add and drop channels at intermediate nodes while allowing pass-through of other channels and the flexibility to add dispersion compensators at each repeater site to minimize transmission penalties. With these technologies, 10 Gb/s long-haul systems have extended their reach well past 1000 km, and 10 Gb/s metropolitan systems have surpassed 300 km in ring circumference.

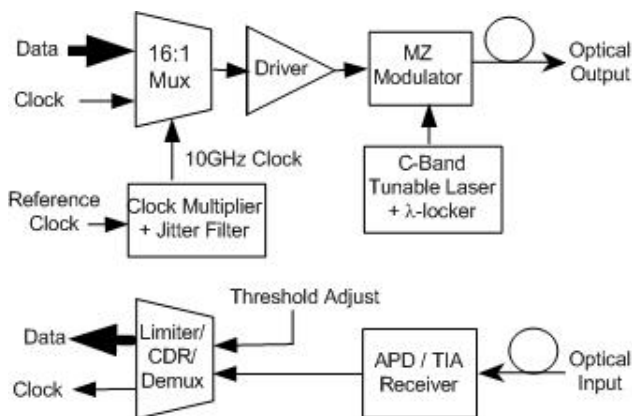


Figure 12: Functional block diagram of a DWDM transceiver

The success of DWDM systems has resulted in a number of technical and logistical challenges. Several system design strategies can mitigate the problems of DWDM transport. Specifically, for a DWDM transceiver, an effective design includes the use of tunable lasers, advanced electronic control circuitry, and high-performance modulators and receivers. Figure 12 depicts a simplified design architecture of a DWDM transceiver.

Tunable Lasers for Channel Sparring

Traditionally, each 10 Gb/s optical channel in a DWDM system is supplied by a unique line card, which specifies

its optical carrier frequency. The large number of channels in the system puts enormous cost pressures on both the telecommunications carrier and the equipment provider when dealing with the inventory and forecasting of line cards. Tunable lasers are useful for reducing the number of inventory codes that an equipment provider must supply when deploying a DWDM system. Full-band tunable lasers can take all the channels in the EDFA bandwidth and reduce them to a common inventory code. Thus, a single tunable line card can be used as a spare for any optical channel in a DWDM system.

One example of a full-band tunable laser is the External Cavity Laser (ECL), which is shown in Figure 13. The LiNbO₃ crystal mirror and the gain-chip output facet, denoted by dark lines in the diagram, provide the optical boundaries (e.g., mirrors or facets) of the laser cavity. Thermally tuning the intra-cavity filters and adjusting the cavity length provide single-mode operation at a chosen channel wavelength. Coupling optics transmit the output to PM fiber and provide a power tap for output power control [11].




Figure 13: Function block diagram of an ECL tunable laser

Channel Monitoring in DWDM Systems

Most carriers require that the DWDM system recognize the presence and absence of channels as they are added and dropped from the network. By applying low-frequency Amplitude Modulated (AM) trace tones to the transceiver output, knowledge of channel presence is established between the terminal equipment and each EDFA repeater site, without degrading other transmission performance parameters. Generally, these requirements fall within a modulation spectrum from 10 kHz to 1 MHz, with peak-to-peak modulation indices ranging from 1% to 5% of the total output power. The AM of tunable lasers is significantly more difficult than that of single wavelength lasers. This is because the DC bias current of a tunable laser varies with each channel and with output power in a complex manner. A sophisticated calibration routine is

required for modulating the tunable laser bias current. One benefit of using tunable lasers such as an ECL is that the gain of the laser can be sufficiently amplitude modulated without producing chirp or frequency modulation.

To account for the above factors, the gain of the transmitter trace circuit for a DWDM transceiver must be adjustable and flat over the AM bandwidth of operation. The circuit also filters any noise that is within the bandwidth of other control circuits, such as the servo signals used for wavelength locking. In Figure 14, the transfer function versus modulation frequency is shown for a tunable DWDM transceiver at a specific channel frequency.

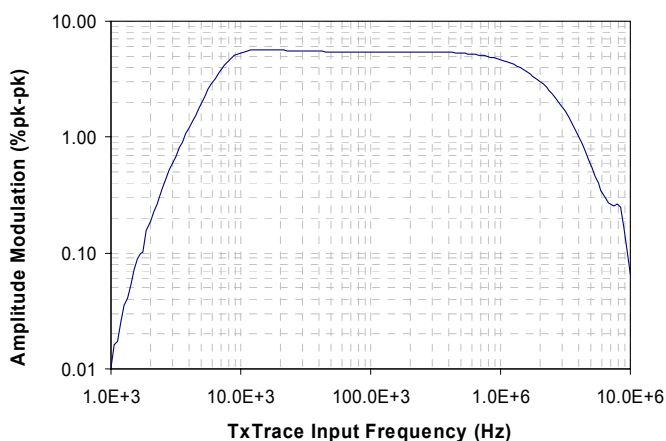


Figure 14: Optical output modulation transfer function (%/mV) as a function of input tone frequency, for an input amplitude of 400 mV_{pk-pk}.

High-Power Launch Conditions

In most DWDM systems, the channel power can be well above 10 dBm for long (> 100 km) fiber spans where intermediate repeater sites are not available. High optical intensities in optical fibers create a number of nonlinearities that cause transmission penalties. One such nonlinearity, which can be mitigated by the transceiver design, is the Stimulated Brillouin Scattering (SBS) of the optical carrier signal.

SBS is a nonlinear phenomenon that coherently beats the optical carrier with the acoustic phonons (i.e., molecular vibrations) of the optical fiber. As a result, the optical channel power is reflected back towards the source. This effect creates bursts of errors at the receiver as the channel power is scattered. Being nonlinear, the amount of SBS increases exponentially with the optical intensity in the fiber. The SBS threshold characterizes the launch power at which the back-scattered light becomes significant. This limits the amount of launch power that

an EDFA can provide to a DWDM system, and thus ultimately limits the transmission distance.

The most common SBS suppression technique effectively broadens the laser linewidth to reduce the power spectral density in the fiber, and thus increases the threshold at which SBS occurs [12]. For high-power applications, linewidth broadening of up to 1 GHz is common. At this linewidth, the EDFA may safely launch more than 19 dBm per channel into standard single-mode fiber.

For tunable external cavity lasers, modulation of the gain medium does not broaden the laser linewidth. This is because the phase of the cavity is dictated mostly by the end mirrors and filters, and less from the index of the gain medium. For ECLs, linewidth broadening is achieved by applying a voltage modulation to the external cavity's LiNbO₃ end mirror, thus producing a pure optical phase modulation [11]. Providing adequate excitation voltage amplitude to the LiNbO₃ crystal is a major circuit design challenge. The step-up transformer is a key component for achieving the necessary amplitude. One of the biggest difficulties regarding the transformer design is achieving high output voltage in a very small size.

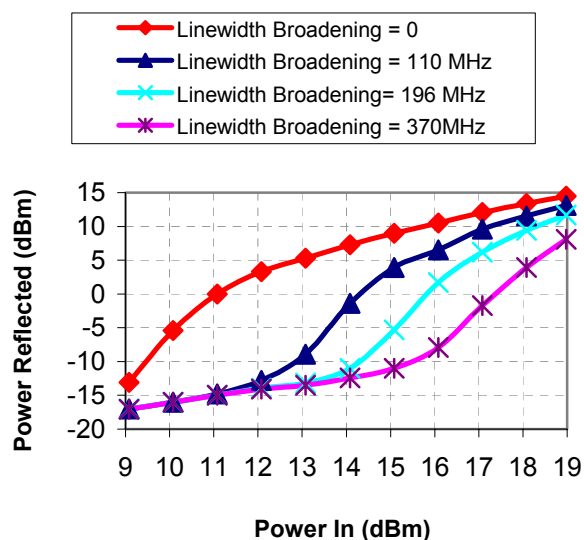


Figure 15: Measured SBS reflected power versus launch power over 100 km of single-mode fiber. The SBS suppression scheme has increased the SBS threshold by more than 7 dB.

By applying direct FM modulation in the tunable laser, the SBS threshold of a DWDM transceiver can be increased. Figure 15 illustrates the measured power reflected back to the source as a function of input power into single-mode fiber. As the linewidth is broadened, less power is reflected. At 15 dBm of input power, for example, we observe, the reflected power drops from about 9 dBm to -12 dBm as the voltage of the end mirror

is applied. This results in an overall 7 dB increase in the SBS threshold. Meanwhile, the resultant AM due to the tone is less than a tenth of a percent of the overall output power.

DWDM Transmission and Optical Noise

The transmission performance of a communications link is typically characterized by its system BER, which is defined as the ratio of errored bits to the total number of bits transmitted. The BER is equal to the complementary error function, $erfc$, of the quality factor, Q , where

$$Q = V_{sig} / \sigma = (\mu_1 - \mu_0) / (\sigma_1 + \sigma_0), \quad (1)$$

and V_{sig} and σ are the received signal and noise levels, respectively [13]. The complementary error function decreases logarithmically as Q increases. In an optical transceiver, the received signal is proportional to the peak-to-peak voltage across the input of the CDR unit, as shown in Figure 15, and we define μ_1 and μ_0 as the mean voltages of the one and zero bits, respectively. The received noise level consists of the mean (rms) noise voltages of the one and zero bits, σ_1 and σ_0 , respectively. In thermally noise limited systems, σ_1 and σ_0 are equal in value. This is not true in amplified DWDM systems.

For optical communication systems, the system BER is a function of the transmitter extinction ratio, the transmitted eye quality, and the receiver tolerance to noise sources. The BER is also a function of several system parameters, such as the input power to the receiver, the amount of dispersion penalty in the fiber, the nonlinear penalty due to high launch powers, and the Optical Signal to Noise Ratio (OSNR). Both the transmitter and receiver portions of a DWDM transceiver can be optimized to improve system BER.




Figure 16: BER vs. received optical power for back to back and after 100 km of single-mode fiber. The data rate is 10.7 Gb/s. The receiver sensitivity is about -28 dBm at a BER of 1E-12.

DWDM Transmitter Design

Equation (1) shows that the Q can be optimized by increasing the transmitter extinction ratio, which is defined by the ratio of the mean one and zero bits. In addition, the Q increases with received signal power. After transmission through fiber, the received signal levels may decrease due to fiber loss and dispersion penalty. For this reason, a high-performance transmitter must output a high optical power and yield a high extinction ratio. As shown in Figure 12, a typical DWDM transceiver employs a high-power tunable laser with a low-loss external LiNbO₃ Mach-Zehnder Modulator (MZM). This combination can result in an output power of greater than 6 dBm.

MZM modulators provide much higher extinction ratios than the electroabsorption modulators that are commonly used in shorter reach applications. Also, they have a well-defined chirp. By properly controlling the modulator chirp, longer transmission distances may be achieved without dispersion compensation. Figure 16 shows the BER as a function of received optical power for a negatively chirped DWDM transceiver before and after 100 km of single-mode fiber. The high transmitter extinction ratio produces an extremely sensitive (-28 dBm) recovery of data using an APD receiver, and the well-controlled chirp only causes 1 dB of dispersion penalty after fiber. The total link loss budget for this DWDM transceiver is more than 33 dB, assuming a 6 dBm launch power.

Optical Noise

In amplified DWDM systems, the loss from the fiber span does not limit the system BER performance because the gain from an EDFA equalizes the fiber loss. In other words, the V_{sig} into the receiver can be maintained constant as long as amplification is available. As a result, amplified DWDM systems are typically noise limited due to the accumulation of optical noise after a chain of EDFAs. This noise is present at the receiver, and it is often referred to as Amplified Spontaneous Emission (ASE) noise. DWDM system designers characterize this noise in terms of an OSNR, where power of the signal in a channel is compared to the optical noise floor. The design of a DWDM transceiver must be able to adapt and optimize the system to various OSNR levels. An approximate expression for OSNR at the end of an optically amplified link with N_{amp} fiber spans preceding each amplifier is given by

$$OSNR (dB) = 58 + P_{out} - L_{span} - NF - 10 \log(N_{amp}) \quad (2)$$

where P_{out} is launch power per channel in dBm, L_{span} is fiber loss in dB of the fiber span, and NF is the EDFA noise figure in dB [14].

At the receiver, the noise from a DWDM system is given by the statistical sum of the electrical and optical noise sources. The electrical noise sources within a receiver itself is composed of dark current, shot noise, and thermal noise. The noise levels for the one and zero bits can be described as

$$\sigma_1^2 = \sigma_{dark}^2 + \sigma_{shot}^2 + \sigma_{thermal}^2 + \sigma_{ASE}^2 + \sigma_{sig-ASE}^2 \quad (3)$$

$$\sigma_0^2 = \sigma_{dark}^2 + \sigma_{shot}^2 + \sigma_{thermal}^2 + \sigma_{ASE}^2 \quad (4)$$

where σ_{dark} is the dark current noise; σ_{shot} is the shot noise; $\sigma_{thermal}$ is the thermal noise; σ_{ASE} is the noise due to ASE from the EDFAs; and $\sigma_{sig-ASE}$ is the noise due the beating between the ASE and the optical signal.

DWDM Receiver Design

When the noise levels are different on the one and zero bits, a more detailed derivation for the system BER is needed [15]. This is given as

$$BER(V) = (1/2) \operatorname{erfc} [(\mu_1 - V) / \sigma_1] + (1/2) \operatorname{erfc} [(V - \mu_0) / \sigma_0], \quad (5)$$

where V is the receiver decision threshold voltage. The maximum system performance is achieved when the BER is minimized, and this occurs when the decision threshold is at its optimum, or

$$V_{opt} = (\sigma_0 \mu_1 + \sigma_1 \mu_0) / (\sigma_1 + \sigma_0). \quad (6)$$

In unamplified systems, the dominant noise term in the receive path is receiver thermal noise because ASE noise is nonexistent. According to Equations (3-4), this implies that the noise is symmetric on the ones and zeroes, and the optimum decision threshold is halfway between the ones and zeroes level (e.g., fifty percent of the signal amplitude) using Equation (6).

The optimum threshold will change in amplified systems where the OSNR is low (about 10 dB) and the ASE noise is dominant. This primarily occurs in long-haul DWDM systems where many EDFAs are used to maximize the transmission distance. As shown in Equations (3-4), the ASE noise from the EDFA is present on both the one and zero bits. However, the noise on the one bit consists of an additional noise term due to beating between the signal and ASE ($\sigma_{sig-ASE}$). This leads to an asymmetry in the noise levels on the one and zero bits. Because the noise on the ones level is larger than that on the zero, the optimum decision threshold level is skewed toward the zero level, as expected, by Equation (6). Note that the $\sigma_{sig-ASE}$ beating noise is also dependent upon the received signal power.

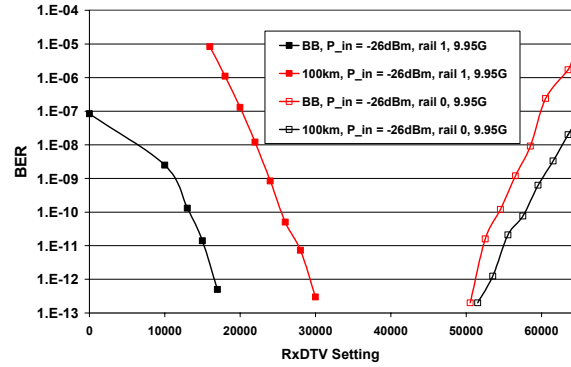


Figure 17: System BER vs. decision threshold setting

For these reasons, long-haul DWDM systems actively control the receiver decision threshold level to minimize the recovered BER under all possible conditions. DWDM receivers are equipped with an external input for adjusting the decision threshold level before the CDR, as shown in Figure 12. An example of the system BER as a function of receiver decision threshold voltage is shown in Figure 17. The BER is minimized at the optimum voltage (near the middle of the range), and it increases as the voltage approaches both the one and zero levels (towards the ends of the range). It is evident that the BER varies asymmetrically as the decision voltage approaches the one and zero levels. This is due to the different noise elements for each level. It is also noted that the threshold moves closer to the zero level after transmission through 100 km of fiber. This is primarily due to chromatic dispersion and the resultant interference on the ones level.

The tolerance of a receiver to optical noise at low system OSNR values is a key figure of merit. This is usually measured at the optimum decision threshold. In Figure 18, the optimum BER for each OSNR level is plotted for a DWDM transceiver with and without transmission through fiber. As shown, the transceiver can recover a BER of 1E-12 in a DWDM system that provides only 17 dB of OSNR. At higher BER levels, such as 6E-5, the system OSNR can be as low as 12 dB. DWDM systems using Forward Error Correction (FEC) will typically operate at these high BER levels.

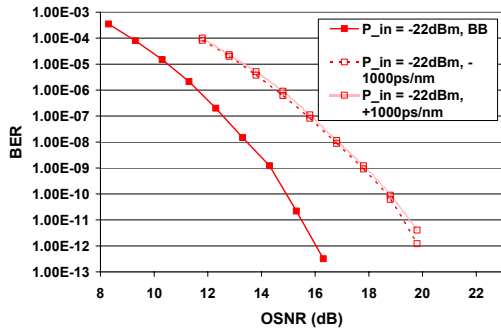


Figure 18: BER performance versus OSNR for a typical tunable DWDM transceiver operating at 11.3 Gb/s. As the system OSNR increases, the BER improves (decreases). The receiver decision threshold is optimized at each condition.

Summary of DWDM Transceivers

The design of DWDM line-side optics is primarily driven by stringent system requirements that have been developed over the last several years based upon the limitations of optically amplified transmission. These optics must be able to satisfy the basic functionality for DWDM metropolitan and long-haul systems, including the ability to provide excellent transmitter extinction ratio, high launch power, and optical noise tolerance. These transceivers must also have several additional features in order to provide a high quality of service to the telecommunications carrier. These include the use of tunable lasers, channel trace tones, and adjustable receiver thresholds. By combining these attributes into a standardized modular interface, the DWDM transceiver will be able to penetrate an existing proprietary market space.

ELECTRONIC DISPERSION COMPENSATION

We now begin a discussion of a technology that can apply to several variations on the basic 10 Gb/s transceivers discussed thus far. Electronic Dispersion Compensation (EDC) is a technology that can be realized in an integrated circuit and thus implemented as an enhancement to various transceivers, enhancing the performance of existing designs and enabling new applications.

Although optical communication is among the fastest high-capacity data transport means available, there are various impairments that limit the capacity per fiber further than the capabilities of the active devices transmitting and receiving light at the ends of the fiber. Besides simple attenuation of the transmitted signal,

various dispersion phenomena affect the ability to recover the signal because of Intersymbol Interference (ISI). In general all dispersion phenomena lead to optical pulse broadening and ISI through different optical properties of the fiber. Recent improvements in the speed of silicon enable deployment of electronic filtering techniques, widely used in lower speed systems such as disk drives and wireless applications. The filter technology applied to optical impairments is known as EDC.

Fiber Dispersion

The most common types of optical dispersion are shown in Figure 19. Chromatic Dispersion (CD) is significant for 10 Gb/s transmission at 1550 nm and distances above 40 km. The modulated spectrum (spread) of the light carrying information and the fact that CD causes light at different wavelengths to propagate at different speeds in fibers causes a broadening of the transmitted light pulses when transmitted over longer distance. The typical means to overcome CD in the optical domain are by reducing unwanted spectrum spread of the transmitter, by inserting sections of fiber with inverse dispersion behavior compared to the standard fiber, or by manipulating the index profile and hence the dispersion coefficient (ps/nm/km) of the fibers. However, all these means are costly and attractive only in long-haul systems. Due to the rather static nature of CD, optical compensation techniques are mostly stationary, set once during installation.

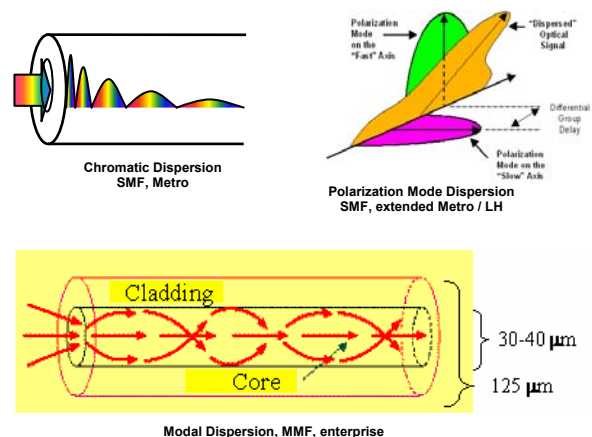


Figure 19: The most important optical dispersion impairments and their relation to markets/applications

Polarization Mode Dispersion (PMD) is a phenomenon occurring from the propagation of polarized light through fiber or optical components. In a polarized environment, the light traveling down the fiber can be split into two orthogonal directions or axes. If the fiber geometry is not exactly circular the light in the two axes will not travel at equal speed (birefringence), and the transmitted pulse will be broadened at the receiver end of the fiber. PMD is

secondary to CD and of less importance in today's fibers due to better control of the geometries in the manufacturing process. This reduces the importance to long-haul systems where CD is compensated in the optical domain; however, PMD can be significant in legacy single-mode fibers as they are upgraded from 2.5 Gb/s to 10 Gb/s transmission. PMD is problematic since it is difficult to compensate for due to its mechanically sensitive nature. Bending the fiber, vibrations, and thermal expansions will all change the actual PMD, and therefore compensation systems have to be adaptive.

Modal dispersion is significant in Multimode Fibers (MMF), widely used in shorter reach applications up to a few hundred meters. It is caused by the fact that the optical pulse splits into several paths (modes) traveling at different speeds down the fiber. The problem is somewhat identical to multi-path fading of wireless signals due to echoes from buildings etc. The use of MMF in enterprise systems has been traditionally related to the cost of connectorizing cables. When dealing with short distances the cost of adding connectors to cables, patch cords, etc. dominates, so there are significant savings in using MMF since the splicing and connectorization processes are more robust and tolerant to misalignments. The availability of low-cost 850 nm multimode VCSEL sources has also lowered the effective cost of multimode links. The penalty paid is modal dispersion, which is significant at speeds beyond 4 Gb/s.

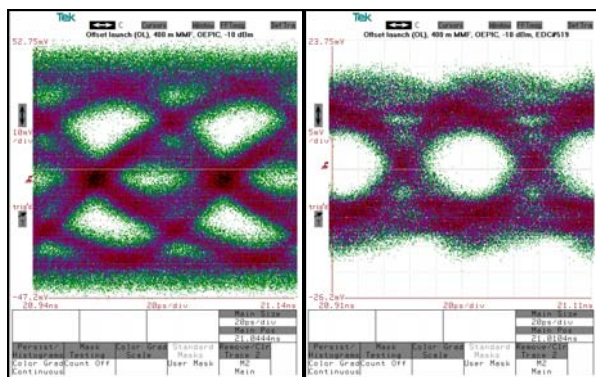


Figure 20: Offset launch eye diagram after 400 m MMF (left), equalized eye diagram (right)

Although the types of dispersion described above can be compensated or controlled in the optical domain, they can with some approximation also be dealt with in the electrical domain. When optical pulses are received by a photodetector, the signal is squared, and phase information of the transmitted signal is lost, leaving only a time continuous power distribution of the received optical pulses.

Applying techniques developed for recovery of information from disk drives, wired communication

systems such as DSL, and wireless systems, the signal can be recovered from a signal that looks totally distorted at a first glance, as shown in Figure 20. This technology, when applied to optical dispersion, is known as EDC.

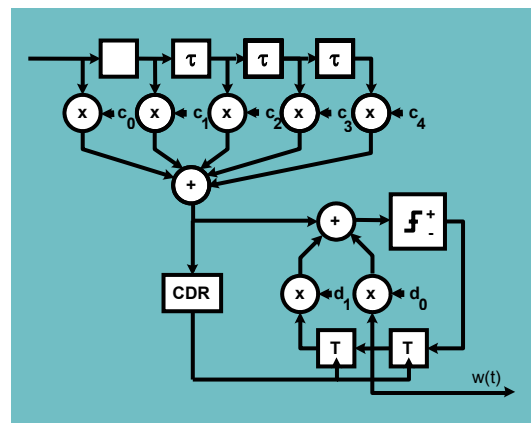


Figure 21: Conceptual block diagram of an electronic dispersion compensation solution, consisting of an FFE, FBE/DFE, and CDR

Adaptive filter techniques are widely described in the literature and can be implemented in a variety of ways. The Finite Impulse Response (FIR) method restores the original signal by reorganizing the received signal from a finite number of time-dispersed weighted partitions that, when correctly adjusted to perform the inverse function of the dispersive channel, will recover the transmitted pulse. The filter consists of a combination of feed forward taps providing linear filtering and feed back taps dealing with the non-linear portion of the pulse reshaping, as shown in Figure 21. Choosing an FIR filtering technique allows the use of adaptive algorithms used to automatically align the filter coefficients to minimize the error relative to an ideal signal. Various adaptation algorithms are possible, among them the Least Mean Square (LMS) approach, which continually compares the filtered output to that of an ideal signal and adjusts the filter coefficients to minimize the square of the error.

Although silicon technology speeds are exceeding the base requirement to operate a 10 Gb/s optical link, carefully architected solutions are needed to best optimize for the lowest possible power consumption. Choosing the right architecture with the optimum balance of FFE and FBE taps to form the optimal combination for the application is critical. Architecting the adaptation control function with the right set of characteristics is an important step to provide the EDC function with the right profile for the given task.




Figure 22: Optical module examples showing the physical content of an EDC-based and a 4-wavelength-based module for 10 Gb/s Ethernet

EDC for Modal Dispersion

The enterprise market is one area of optical transmission that will greatly benefit from the use of filtering techniques. This market is typically Ethernet oriented, and the increments of speed from one generation to the next are one order of magnitude, 10 Mb/s – 100 Mb/s – 1 Gb/s – 10 Gb/s. The market is cost-sensitive and requires robust and reliable solutions that operate under office environment conditions with plug-and-play capabilities. The use of MMF is widespread from the early introduction of 1 Gigabit Ethernet (GbE). Because they offered a good combination of longer distance and higher speed enabled by optical transmission, these fibers were installed as backbone connections in networks, vertical risers in large buildings, campuses and other applications spanning up to 300 m distances. As desktop and server connections increase in speed from 100 Mb/s to 1 Gb/s, interest has increased in upgrading backbone fiber links from 1 GbE to 10 GbE.

Upgrading existing fibers originally installed for 1 GbE to run 10 GbE is not easily done, since the modal dispersion is significant and unless properly compensated, could prevent error-free transmission at 10 Gb/s. The current solution to overcome modal dispersion at 10 Gb/s in 300 m legacy multimode links is to multiplex 4 individual data streams at ¼ of the datarate onto a fiber at individual optical wavelengths. However, this solution is relatively expensive and is not scalable to smaller form factors (notably XFP). Electronic dispersion compensation offers a cost-effective alternative. As there are no added system blocks, over time the EDC feature will be integrated with the PCS block as shown in Figure 22.

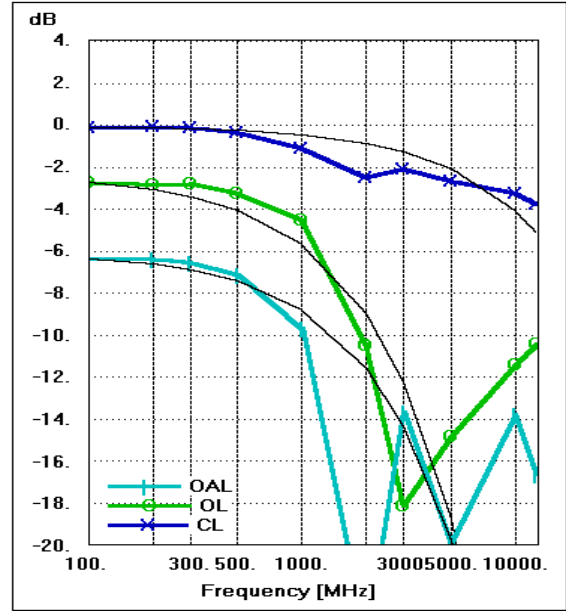


Figure 23: Modal bandwidth measurement of 400 m of low BW multimode fiber under different launch conditions, center launch, offset launch, and overfilled launch with added modes

IEEE Task Force, 10GBASE-LRM

A solution based on established filter theory and implemented in silicon is being pursued by development teams throughout the industry. The IEEE has recently established a task force, 10GBASE-LRM, working to standardize the application space defined by a channel description and a required minimum probability of correctly compensating a worst-case fiber based on statistically collected data describing the base of installed fibers in today’s network. This is not a trivial task since modal dispersion is very dependent on the mechanical and optical characteristics of the fiber, including the index profile, where fibers installed in the early 1990’s exhibit worse behavior than new fibers. So far, fibers have been characterized using a simple bandwidth measurement, but when dealing with dispersion compensation this is an inaccurate method. Two fibers exhibiting the same modal bandwidth may exhibit very different dispersion patterns of which one could be significantly more complex to compensate than the other.

In Figure 23 the three measured curves show the modal bandwidth of one fiber under different transmitter launch conditions (the way light is coupled into the transmit end of the fiber). Launch conditions are important characteristics that are also covered by the IEEE 10GBASE-LRM task force, which will eventually specify launch condition requirements for 10 Gb/s transmission over multimode fibers.

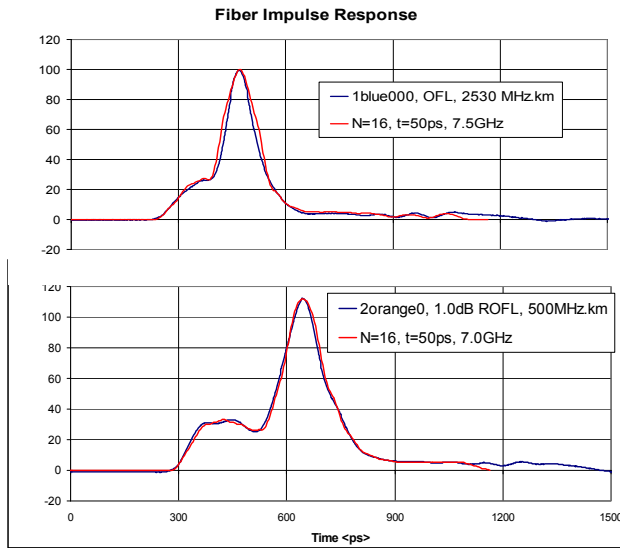


Figure 24: Example of realistic impulse response from multimode fibers of 300 m length

A better way to characterize multimode fibers is by means of its impulse response or by recording a bit pattern from which the pulse response can be obtained, as shown in Figure 24. Once the impulse response of the channel (the fiber with a given launch condition) is known, it is a simple mathematical task to compute the required filter combination that exhibits the inverse function of the fiber and therefore would compensate the dispersion.

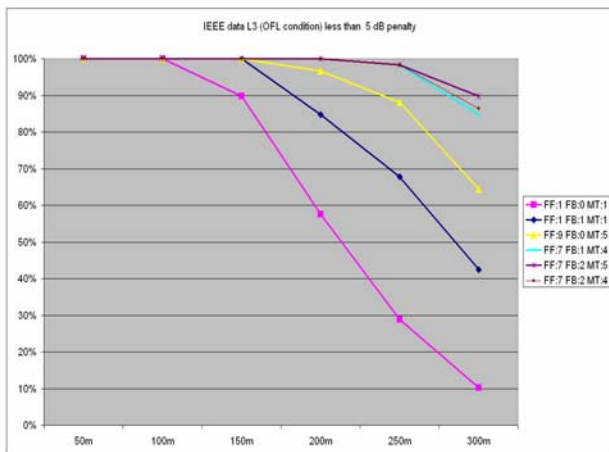


Figure 25: Analysis of a selected set of worst-case fibers and the ability to compensate the modal dispersion at a optical power penalty of 5 dB. The individual curves plot particular filter combinations, with the number of FFE taps and FBE taps listed.

Because the requirement to compensate modal dispersion is related to all existing fibers installed, there are requirements to cover a sufficiently large portion of these fibers with the chosen adaptive solution. This allows the

EDC to adapt to any fiber within a defined set of characterized fibers, representing all existing fibers. Extensive modeling of filter combinations and their ability to restore signals from representative worst-case fibers are used in architecting the required solution. In Figure 25 a set of curves, each representing a specific filter combination, shows the percentage of fiber responses successfully compensated to a BER of 10^{-12} at a given optical power penalty.

Test Circuit Development

Although the raw speed of silicon technology has increased and allows more signal-processing circuit techniques in 10 Gb/s systems, there is no room for complex techniques to compensate dispersion in enterprise applications. This is because sophisticated filtering would require strong parallel approaches to provide the required performance, with excess power consumption that is incompatible with existing power budgets for optical modules. The challenge therefore lies in the circuit designer’s ability to provide circuits fulfilling the requirements of the defined channels while at the same time accommodating the power budget by integrating the EDC feature into a transceiver building block in the enterprise optical transceiver module. As pictured in Figure 21, the block can be described using filter block terminology such as number of taps, tap weight range, feed forward equalizer, feed back equalizer, etc. The individual block characteristics and tradeoffs are all measured in conventional metrics such as linearity, bandwidth, resolution, and sensitivity. Breaking down the circuit into these well-defined equivalent blocks help the circuit designer to choose the optimal filter combination required for the application while at the same time optimizing the circuit.

Measurements from experimental circuits exhibit predicted performance improvement when applied to modal dispersion. Figure 26 shows the performance improvement at different launch conditions of a 400 m fiber link using a test device with a 5-tap FFE. The graph plots BER vs. receiver Optical Input Power (ROP), a common way to express the performance of optical systems. With Center Launch (CL) conditions the modal dispersion is limited, which explains the low relative improvement. With Offset Launch (OL) the eye diagram of the received signal and the equalized signal are those of Figure 20 resulting in a 4.2 dB improvement in ROP at BER = 10^{-3} and a dramatically higher improvement at lower BER (not measured due to limitations in setup). From the diagram, it can also be seen that the residual dispersion penalty (for offset launch) is 2.5 dB compared to the back-to-back reference measurement (transmitter coupled directly to receiver with 1-2 m fiber).

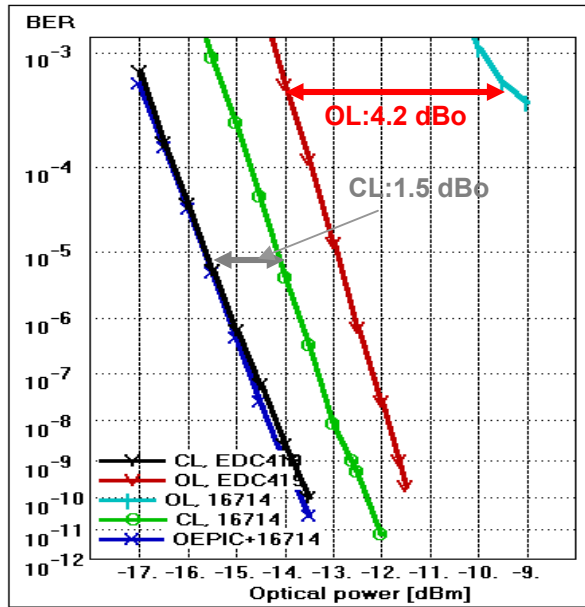


Figure 26: Optical penalty comparisons to unequalized link. Center launch improvement, 1.5 dB; offset launch improvement, 4.2 dB (400 m MMF).

Intel has had a long-standing interest in adaptive filter technology and has benefited from a strong knowledge pool derived in dealing with DSL and wireless applications. This has been used in the technology development providing fundamental building blocks for EDC to be integrated into products directed at the enterprise and metro market segments. The proven technology has demonstrated superior performance and helped set the expectations in the industry. Intel has contributed to the work on multimode EDC in the IEEE standards effort and has provided strong data, based on both modeling and experimental work, paving the way for an industry-wide acceptance of EDC-enabled transmission of 10 Gb Ethernet over legacy multimode fibers.

CONCLUSION

The state of the art in 10 Gb/s optical transceiver design has been presented, with the architecture, functionality, and performance criteria for each of the major functional elements. Based on this basic structure, many different standardized products are available, with various form factors, fiber reaches, performance levels, and costs to satisfy the needs of optical communication system-level designs. Two emerging technologies that are pushing performance to new levels while reducing the cost of implementing modern designs were also explored. The wavelength tunable transceiver allows system vendors who had previously built DWDM solutions at the board

level to use modularized transceivers in their systems. The promising technology of electronic dispersion compensation is expanding link budgets and margins while driving new thinking about the capability of fiber systems. While the application to compensate for modal dispersion in multimode fiber links was discussed here, the technology lends itself to a variety of systems where dispersion and bandwidth limitations are present. As bandwidth demands continue to increase and 10 Gb/s links are adopted more widely in the industry, technologies such as these will play a significant role in satisfying that demand.

ACKNOWLEDGMENTS

We acknowledge the hard work and innovations of the engineers of Intel's Optical Platform Division (OPD) and Optical Networking Components Division (OND). Special thanks to Scott Schube for insights and editing. Much of the work in this paper would not be possible without the knowledge, insight, and patience of Steve Buchheit.

REFERENCES

- [1] 300 Pin MSA Group, *300 Pin MSA Edition 4*, August 2002, available: www.300pinmsa.org.
- [2] XENPAK MSA Group, *XENPAK MSA Issue 3.0*, September 2002, available: <http://www.xenpak.org>.
- [3] X2 MSA Group, *X2 MSA Issue 1.0b*, February 2003, available: <http://www.x2msa.org>.
- [4] XPAK MSA Group, *XPAK MSA v2.3*, August 2003, available: <http://www.xpak.org>.
- [5] 10-Gigabit Small Form-factor Pluggable (XFP) Module Group, *XFP MSA Revision 3.1*, April 2003, available: <http://www.xfpmsa.org>.
- [6] Optical Internetworking Forum, *SFI-4 (OC-192 Serdes-Framer Interface) OIF-PLL-02.0*, "Proposal for a common electrical interface between SONET framer and serializer/deserializer parts for OC-192 interfaces," Fremont, CA, September 2000.
- [7] International Telecommunication Union, *ITU-T Recommendation G.691*, "Optical interfaces for single channel STM-64, STM-256 and other SDH systems with optical amplifiers," Montreal, Quebec, Canada, October 2000.
- [8] Telcordia Technologies, *Synchronous Optical Network (SONET) Transport Systems*, "Common Generic Criteria (GR-253-CORE)," Piscataway, New Jersey, September 2000.
- [9] IEEE Standard 802.3ae, "IEEE Standard for Carrier Sense Multiple Access with Collision Detection

(CSMA/CD) Access Method and Physical Layer Specifications-Media Access Control (MAC) Parameters, Physical Layer and Management Parameters for 10 Gb/s Operation,” 2002, available: <http://standards.ieee.org/getieee802/802.3.html>

- [10] Telcordia Technologies, “GR-468-CORE, Issue 1, Generic Reliability Assurance Requirements for Optoelectronic Devices Used In Telecommunications Equipment,” December 1998.
- [11] M. Finot, M. McDonald, A. Daiber, W. Chapman, D. Li, M. Epitoux, E. Zbinden, J. Bennett, W. Kozlovsky, J.M. Verdiell, “Automated Optical Packaging Technology for 10Gb/s Transceivers and its Application to a Low-Cost Full C-Band Tunable Transmitter,” *Intel Technology Journal*, Volume 8, Issue 2, May 2004.
- [12] R. Ramaswami and K. Sivarajan, *Optical Networks-A Practical Perspective*, Academic Press, San Francisco, CA 1998, pp. 238-240.
- [13] N. Bergano, F. W. Kerfoot, and C. R. Davidson, “Margin measurements in optical amplifier systems,” *IEEE Photonics Technology Letters*, vol. 5, no. 3, pp. 304-306, 1993.
- [14] J. L. Zyskind, J. Nagel, and H. D. Kidorf, “Erbium-Doped Fiber Amplifiers for Optical Communications,” in *Optical Fiber Communications IIIB*, I.P. Kaminow and T.L. Koch, eds., pp. 13-68, Academic Press, San Diego, 1997.
- [15] J. M. Wozencraft and I. M. Jacobs, *Principles of Communication Engineering*, Wiley, New York, 1965, pp. 77-84.

AUTHORS' BIOGRAPHIES

Peter Kirkpatrick received his M.S.E.E. degree from the University of Colorado at Boulder, with an emphasis on Microwave Antennas, Circuits, and Optics, in 1999. He worked at the National Institute of Standards and Technology (NIST) and was an early contributor to LightLogic. His current position is senior R&D engineer in advanced development with Intel's Communication Infrastructure Group. He has published articles and has patents pending in the areas of electrical transport, microwaves, RF photonics, and packaging technologies. His e-mail is pete.kirkpatrick@intel.com.

Wei-chiao Fang received his Ph. D. degree in Electrical and Computer Engineering from the University of Illinois in 1997. Dr. Fang then joined the Optical Networking Group of Bell Laboratories, Lucent Technologies. From 2000 to 2003, Dr. Fang headed the development of optical technologies for Lucent's DWDM Metro and Long-Haul business units. He now leads the DWDM

transponder engineering group at Intel Corporation. Dr. Fang has filed several patents in the areas of optical monitoring and 10G modulation techniques. His e-mail is wei-chiao.w.fang@intel.com.

Henrik Johansen attended the Engineering High School of Odense, Denmark, 1979 with an emphasis on Industrial Automation and Servo Systems. He is currently a principal engineer at Intel's Communications Infrastructure Group's Optical Networking Division where he is focusing on high-speed optoelectronic devices. Since 1989 Henrik was an early contributor to GIGA A/S and participated in many worlds' firsts in the 2,5 Gb/s and 10 Gb/s Serdes area until GIGA was acquired by Intel in 2000. He holds three international patents and has several in progress. His e-mail is henrik.johansen@intel.com.

Dr. Benny Christensen graduated from the Technical University of Denmark in 1985, with an emphasis on RF Electrical Engineering and Servo Systems. He earned his Ph.D. degree within the area of Optical Soliton Communications at Tele Denmark Research. Benny was employed at GIGA A/S and has participated in the development of the 2,5 Gb/s and 10 Gb/s Serdes devices. He is currently a senior system and applications engineer at Intel's Optical Networking Division working on high-speed PHY and optoelectronic devices. He holds two international patents and has several in progress. His e-mail is benny.christensen@intel.com.

Jesper Hanberg graduated in Physics from the University of Copenhagen in 1983 and received his Ph.D. degree from that university in 1987. Since 1988 he has worked with design, fabrication, and integration of high-speed optoelectronic components first in Tele Denmark Research and from 1996 in Giga A/S (acquired by Intel in 2000). Currently he is in Intel's Communications Infrastructure Group's Optical Networking Division where he is working on simulation tools for electronic dispersion compensation. His e-mail is jesper.hanberg@intel.com.

Dr. Martin Lobel graduated from the Technical University of Denmark in 1994 with a thesis on Actively Mode-Locked Fiber Ring Lasers for Optical Communication Systems. He holds an M.S. degree in Electrical and Optical Engineering and a Ph.D. degree in Applied Physics from Risoe National Laboratory in Denmark. He is currently with Standards and Advanced Technology in the Optical Networking Component Division at Intel. Before joining Intel, he worked on laser development for GiGA A/S. Martin is an Intel mentor for several Intel-sponsored research activities at universities in Europe, has published many papers, and holds two patents. His e-mail is martin.lobel@intel.com.

Thomas B. Mader received a Ph.D. degree in Electrical Engineering from the University of Colorado at Boulder in 1995, as well as a B.S. degree in Electrical Engineering from the University of California at Berkeley in 1990. He is currently a principal engineer at Intel Corporation, Optical Platform Division (Newark, CA). He was an early member of LightLogic, Inc., a transponder and optical automated-packaging startup, which was acquired by Intel in 2001. His e-mail is tom.mader at intel.com.

Song Quan Shang received his Ph. D. in Physics from the State University of New York in 1991. He is a senior staff development engineer at Intel's Optical Platform Division. He is the leading system designer for Intel's DWDM transponder. Prior to joining Intel, Dr. Shang was a principle engineer at Ciena Corporation where he had successfully developed many leading-edge long-haul and metro DWDM transport products. He is the author and/or co-author of more than 20 papers in peer-reviewed journals on optical technology and communications. His e-mail is song.q.shang at intel.com.

Craig Schulz received his B.S.E.E. degree from Georgia Institute of Technology in 1984 and his M.S.E.E. degree from the University of Southern California in 1986. Craig joined Hughes Space and Communications Company in 1984 where he worked on satellite communication payloads, designing microwave circuits, and optical and bandwidth efficient communication subsystems. He has worked for four years at Intel Corporation and its LightLogic acquisition where he has developed several 10 Gb/s transponder product lines enabling short reach to long-haul DWDM fiber communications. He has filed several patents in the areas of high rate modulation and optical transmit and receive circuits. Currently Craig is the lead design engineer for the Monarch2 long-haul DWDM tunable transponder. His e-mail is craig.schulz at intel.com.

Doug Sprock earned an M.S.E.E. degree in 1998. During his 15 years of industry experience he has designed and tested many electronics subsystems. In 2000 Doug began designing tunable laser control electronics at New Focus where he became the lead electronics project engineer. He holds several patents. His e-mail is doug.sprock at intel.com.

Dr. Jean-Marc Verdiell is an Intel Fellow and director of Optical Technology. He holds a Ph.D. degree in Optoelectronics from the University of Paris. Dr. Verdiell worked at AT&T Bell Laboratories and SDL, pioneering advanced optoelectronic components. He was the founder and CTO of LightLogic, now Intel's Optical Platform Division, which designs and manufactures 10 Gb/s transponders for communications networks. Dr. Verdiell holds 15 patents and authored over 50 peer-reviewed papers. His e-mail is marc.verdiell at intel.com.

Copyright © Intel Corporation 2004. This publication was downloaded from <http://developer.intel.com/>.

Legal notices at <http://www.intel.com/sites/corporate/tradmarx.htm>.

THIS PAGE INTENTIONALLY LEFT BLANK

For further information visit:

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