Reference Guide on Optical Interconnects For High Performance Compute (HPC) Systems

Prepared by:

Nathan Harff, James Kruchowski, Mark Nelson,

Brian Shamblin, and Vladimir Sokolov

Special Purpose Processor Development Group

Mayo Clinic

Rochester, MN 55905

Phone: (507)284-4056

October 2009

Mayo-R-09-39-R1

Table of Contents

List of Figures

List of Tables

Acknowledgements

The authors thank Terri Funk for editing and formatting the text, Steve Richardson, Elaine Doherty, Terri Funk, and Deanna Jensen for figure preparation and Dan Schwab for reviewing the final draft of this Reference Guide.

1 Basic Theory

This section introduces the concepts and basic theory of optical data communications systems, particularly as they may apply to high performance computing (HPC) systems.

1.1 Generic Electro-Optic Link

The components of a generic electro-optic link are shown in Figure 1-1*(A)*. The light is generated in the light source, which is either a light emitting diode (LED) or a laser, and the data to be transmitted (typically non return to zero (NRZ) digital data) modulates this light in the modulator stage. The modulation can be done directly in the laser, by modulating the laser's drive current, or a separate modulator can be used. Different modulation schemes can be used to encode the data on the light. Next the light is transmitted over the transmission medium, which is generally an optical fiber but can also be free space or another type of optical waveguide such as a polymer waveguide in an optical backplane. Finally, the light is received by the photodetectors and the resulting photodetector current is converted to voltage and amplified in the transimpedance amplifier (TIA). The network diagram of Figure 1-1*(B)* shows a specific example of the use of dense wave division multiplexing (DWDM) and single-mode optical waveguides to broadcast the data from multiple transmitters to all receivers. In this approach, light from different modules, each with different transmitted wavelengths, is combined in the array waveguide grating (AWG) (see Section 1.6.5) onto a single fiber and the star coupler (see Section 1.6.4) distributes this multi-wavelength light to all other modules through the transmission medium. Another AWG is used "in reverse" to separate the different wavelengths of light before detection at each receiver.

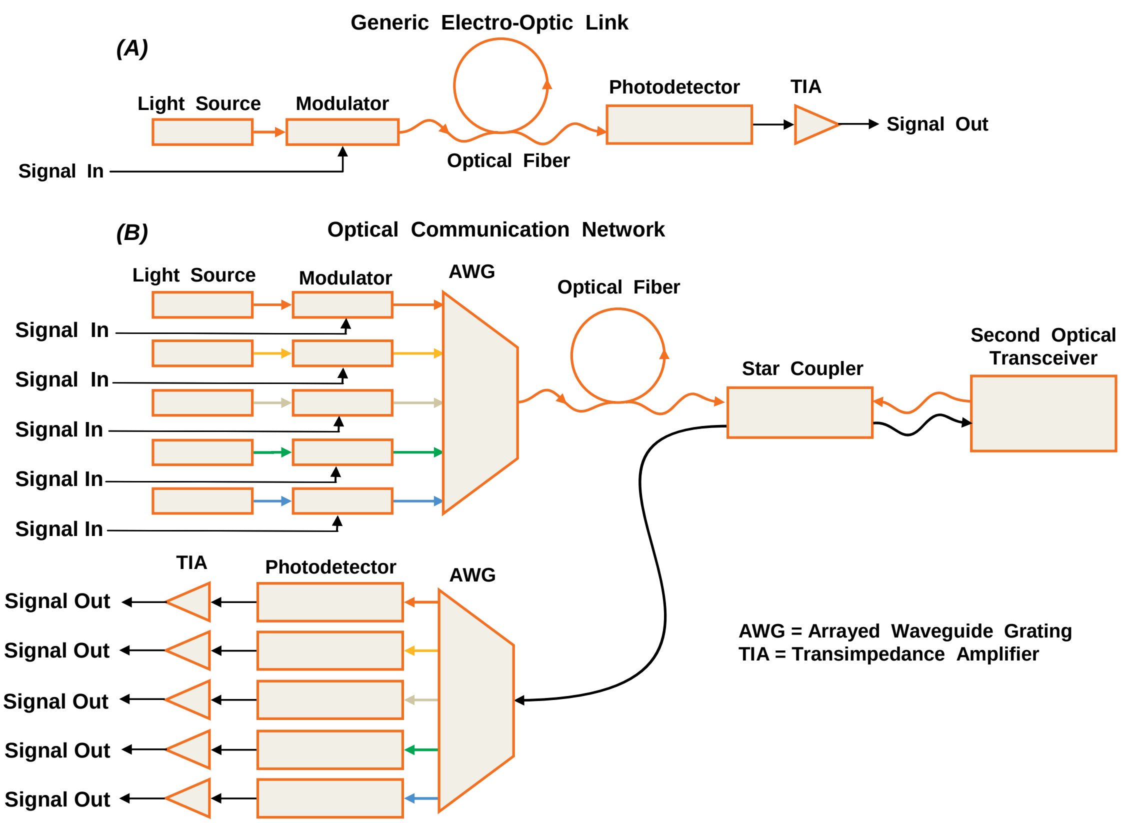

Figure 1-1: Diagrams of *(A)* Generic Electro-Optic Link and *(B)* An Example Communication Network. (40275)

1.2 Guided Light

This subsection briefly discusses the theory behind light transmission in fiber optic communication systems. Section 2 gives details on the common fiber types and their typical characteristics. The requirement for the transmission medium is to get the light from the transmitter to the receiver with acceptably low loss (minimal attenuation), low distortion/dispersion and low crosstalk between different channels. The transmission media can be divided into two categories; free-space and waveguide.

1.2.1 Light Transmission in Free Space Versus an Optical Waveguide

In a free-space medium, a collimated beam of light (typically from a laser) is directed from the source to the detector using mirrors to point the light at the destination. One challenge with free-space systems is the difficulty in aligning the mirrors and keeping them in position so the light reaches the detector. In HPC applications, these systems can work for shorter distances (< ~1m) but are impractical at longer distances due to problems with aligning the beam to the detector and if large numbers of channels are needed.

In a waveguide medium, light is guided from source to detector through the medium. This is accomplished using total internal reflection to confine the light to the guide. As illustrated in Figure 1-2, total internal reflection occurs when light in a material with a higher refractive index (RI, n1) is incident at a large angle (from normal to the interface) on a material with a lower refractive index (n2). In this case, all of the light is reflected back into the waveguide and none of it is lost through transmission into the second material.

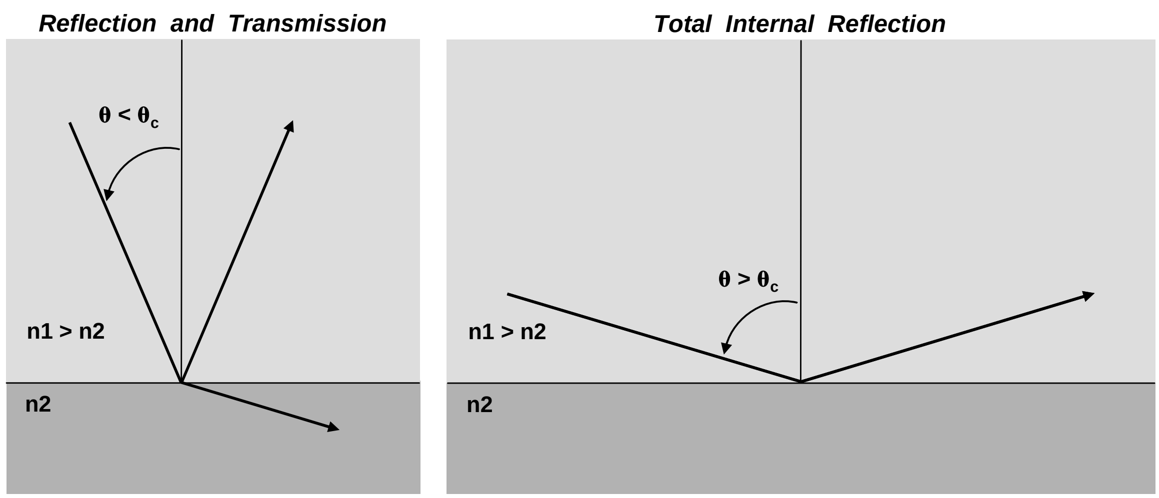

Figure 1-2: Illustration Showing Standard Reflection and Total Internal Reflection. (40276)

1.2.2 Types of Optical Waveguides

Optical waveguides typically have either rectangular or circular cross sections, which are shown in Figure 1-3. Rectangular waveguides are generally fabricated on semiconductors by etching a ridge onto the surface of the wafer or by depositing a dielectric on the semiconductor surface. These are used within lasers and modulators and for guiding light from one device to another in a photonic integrated circuit. Waveguides with rectangular cross sections are also being fabricated on printed circuit boards for chip-to-chip communication on the board. The lower RI material can be either air or a dielectric deposited on the ridge waveguide.

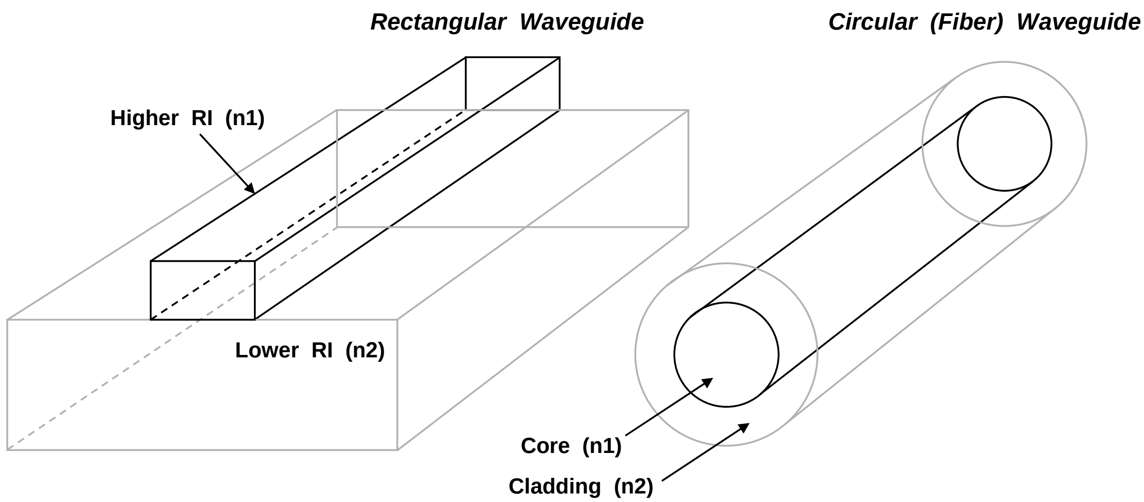

Figure 1-3: Typical Cross Sections of Optical Waveguides. (40277)

The circular waveguides are optical fibers and are typically made of silica glass or plastic depending on the application. They usually have additional outer layers for mechanical protection which are not involved in guiding the light down the fiber. The higher RI center of the fiber is called the core and the outer lower RI material is called the cladding. The difference in the RI between the core and the cladding can be as small as 1% in fibers. The RI profile can be a step index in which the refractive index abruptly changes, it can be a graded index in which it gradually changes or it can have a more complicated profile.

Light being launched into a fiber must be at an angle which allows total internal reflection to occur. If the launch angle (with respect to the fiber axis) is too large, some of the light will be transmitted into the cladding every time the beam reflects off of the cladding and the beam will not reach the other end of a long fiber. The maximum angle of launch that results in total internal reflection defines the acceptance cone of the fiber and the sine of this angle is defined as the numeric aperture (NA) of the fiber. This is shown in the middle section of Figure 1-4.

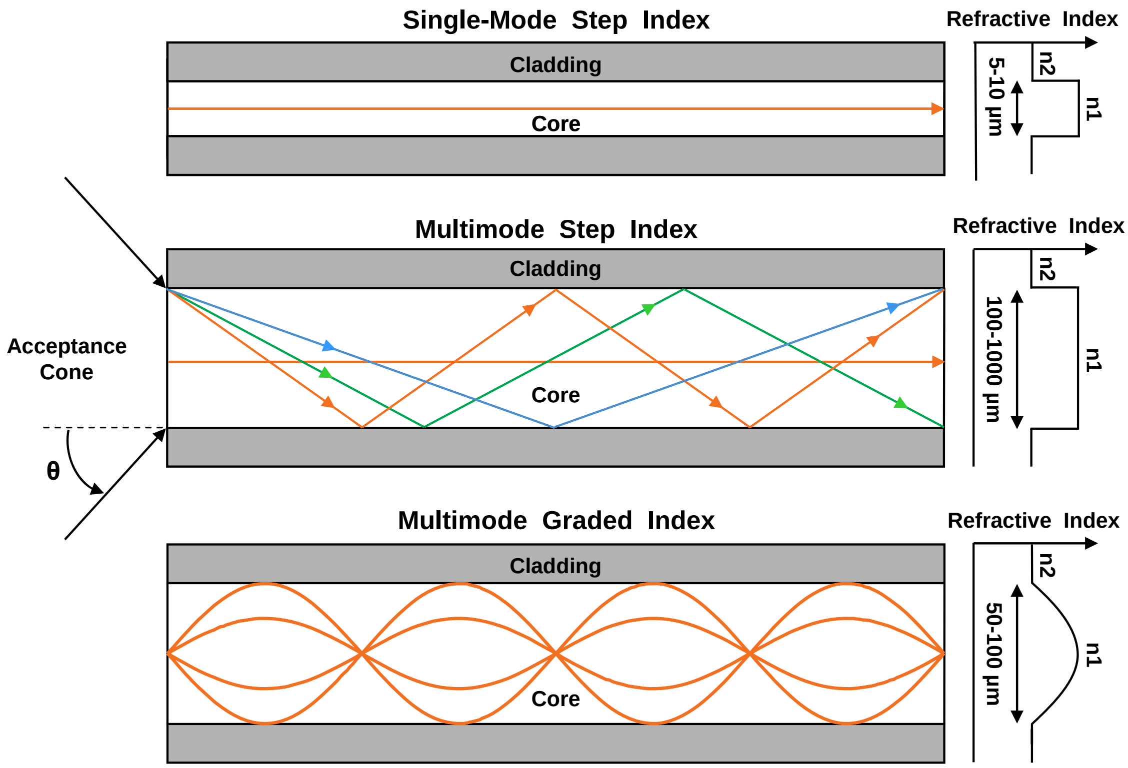

Figure 1-4: Types of Optical Fibers. (40278)

1.2.3 Single-Mode and Multimode Optical Fibers

The light propagating in a fiber can be visualized using ray optics as a beam bouncing off the core-cladding interface on the top of the fiber and then on the bottom and so on from one end of the fiber to another. To fully understand light propagation in fibers it is necessary to consider the solutions of Maxwell's equations for the electro-magnetic fields within the fiber. We will not go into details of solving these equations but will discuss the results. There are multiple solutions to these equations corresponding to different modes, similar to the modes in microwave waveguides. A fiber with a small core (~5-10 µm) will support one mode and is therefore referred to as a single-mode fiber. This mode can be visualized as a single ray propagating down the center of the fiber. A fiber with a larger core (~50-1000 µm) can support multiple modes and is referred to as a multimode fiber. These modes can be visualized as rays bouncing off the core-cladding interface with each mode at a different angle. The larger core also results in a larger NA making it easier to launch light into the fiber.

1.2.4 Attenuation in Optical Fibers

As light travels down the fiber its intensity or amplitude will decrease. Figure 1-5 shows the typical attenuation for single-mode and multimode fibers. There are three main causes for this attenuation. The most significant is Rayleigh scattering which is caused by small variations (< $1/10^{th}$ of the wavelength) in the density and composition of the glass in the fiber that causes the light to scatter. The scattering is inversely proportional to the wavelength to the 4^{th} power so longer wavelengths have lower attenuation. The second attenuation contribution is light absorption by unwanted impurities in the fiber. Pure silica absorbs very little light from ~800 nm to 1700 nm but impurities in the silica do absorb light. With improved purification techniques, the levels of most impurities are insignificant except for water. The oxygen-hydrogen (OH) bond in water absorbs light near 1385 nm causing increased attenuation around this wavelength. There are also smaller peaks (not shown here) in attenuation near 900 nm and 1200 nm. The third attenuation contribution is light absorption by intentional impurities in the fiber. These impurities are generally added to change the RI of the silica. The wavelengths typically used are around 850 nm, 1300 nm, and 1550 nm. The 850 nm band is used because early fibers (1970s and early 1980s) had a minimum in attenuation around this wavelength and it is still used for short-distance communications. The 1300 and 1550 nm bands are currently used because of the low attenuation in these bands.

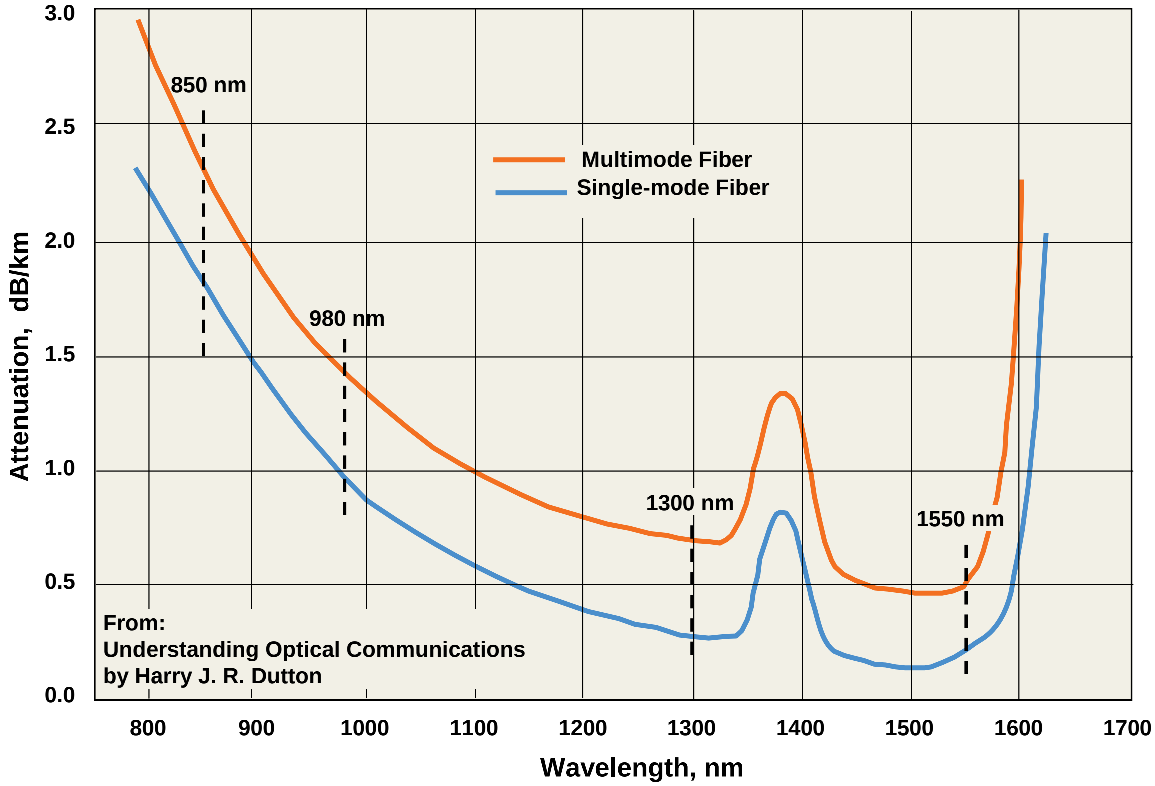

Figure 1-5: Typical Attenuation In Optical Fibers. (40279)

1.2.5 Dispersion in Optical Fibers

Another important characteristic of an optical fiber is dispersion which causes a pulse of light to spread out as it travels along the fiber. If pulses become too broad, they will overlap each other making it impossible to recover the data at the receiver. Dispersion can be put in two classes; modal and chromatic.

Modal dispersion results from different effective velocities for different modes in multimode fiber. The lowest order mode is launched at zero angle and the higher order modes are launched at increasingly higher angles and thus have increasingly more reflections off of the core-cladding interface resulting in longer path lengths which gives them lower effective velocities than modes with fewer reflections. In step-index multimode fibers this results in signal distortion at the far end of the fiber because it takes higher order modes longer to travel the length of the fiber. To reduce this dispersion, graded-index multimode fibers are used because the profile causes light to travel along sinusoidal paths giving the various modes roughly the same effective velocities due to the fact that the phase velocity is inversely proportional to the RI. Light travels slower near the center of the core (highest n) and faster near the core-cladding interface (lowest n). So lower order modes travel along shorter paths near the center of the fiber but have lower phase velocity because of the higher RI. The higher order modes travel along longer paths that extend out to the edge of the core but have higher phase velocity because of the lower RI further out from the core. Modal dispersion is given in units of either MHz•km or their reciprocal in ps/km.

Chromatic dispersion occurs because the velocity of light in the fiber depends on the wavelength of the light. The optical pulses from lasers and LEDs used in optical systems are not at a single wavelength but have a range of wavelengths depending on the source. These different wavelengths travel at different velocities and thus broaden the optical pulse as it travels along the fiber. Chromatic dispersion can be further divided into material dispersion, waveguide dispersion and profile dispersion. Material dispersion results from the wavelength dependence of the fiber's RI. Waveguide dispersion results from the wavelength dependence of the modal characteristics of the fiber. Profile dispersion results from the wavelength dependence of the RI of the core and cladding. Chromatic dispersion is defined as minus the change of the pulse travel time (ps) per unit length of fiber (km) per change of wavelength (nm) and is given in units of ps/(km•nm).

The amount of dispersion increases with the fiber length which limits the useable optical bandwidth for longer fibers. A fiber transmission system can be characterized by the bandwidth-distance product which is typically expressed in units of MHz•km. This indicates how far a signal with a given bandwidth can be expected to travel and still be correctly received at the end of the fiber. Nominally, a signal with twice the bandwidth could travel only half the distance. Strictly speaking the bandwidth-distance product is not totally independent of distance, but practically this figure of merit is nevertheless useful for comparing various link implementations. The required bandwidth is typically 0.5 – 0.8 times the bit rate of the signal.

1.3 Semiconductor Light Sources

1.3.1 Light Emission From a Solid

Light can be emitted from a material when an electron moves from a higher energy state to a lower energy state. The wavelength (λ) of the emitted photon is

$$\lambda\ (\mu m) = h \cdot c/\Delta E = 1.24/\Delta E(eV),$$

where h is Planck's constant (6.63e-34 J-s), c is the speed of light in vacuum (3.00e8 m/s) and ΔE is the energy difference between the two states given in electron volts (1 eV = 1.60e-19 J).

In equilibrium electrons are in the lowest available energy states so they need to be excited to the higher energy state in order for light emission to occur. For lasers, this is referred to as "pumping." Typically lasers are either electrically pumped or optically pumped. In electrical pumping, either a voltage is applied across the material or a current is forced through it to excite the electrons to the higher state. This is discussed in detail in the next section. In optical pumping, electrons are excited to a state higher in energy than the desired state by absorbing light with a shorter wavelength (higher energy) than the desired laser output wavelength. The electrons then relax to the desired state by giving energy in the form of heat to the material. These are known as non-radiative transitions. Once the electrons are in the desired excited state, they emit photons when they move back to the ground state.

1.3.2 Energy Levels and Energy Band Diagrams in Semiconductors

Quantum mechanical theory shows that for individual atoms, there are discrete electron energy levels for the allowed electron orbitals. As atoms are brought together to form a solid, these discrete energy states coalesce into allowed and forbidden energy bands as shown in Figure 1-6. At absolute zero temperature (0 Kelvin = -273.15 C), the lowest energy states are filled by the electrons of the atoms and the higher energy states are empty. At higher temperatures, electrons can move to the higher electron states if they have enough thermal energy to reach an empty higher state.

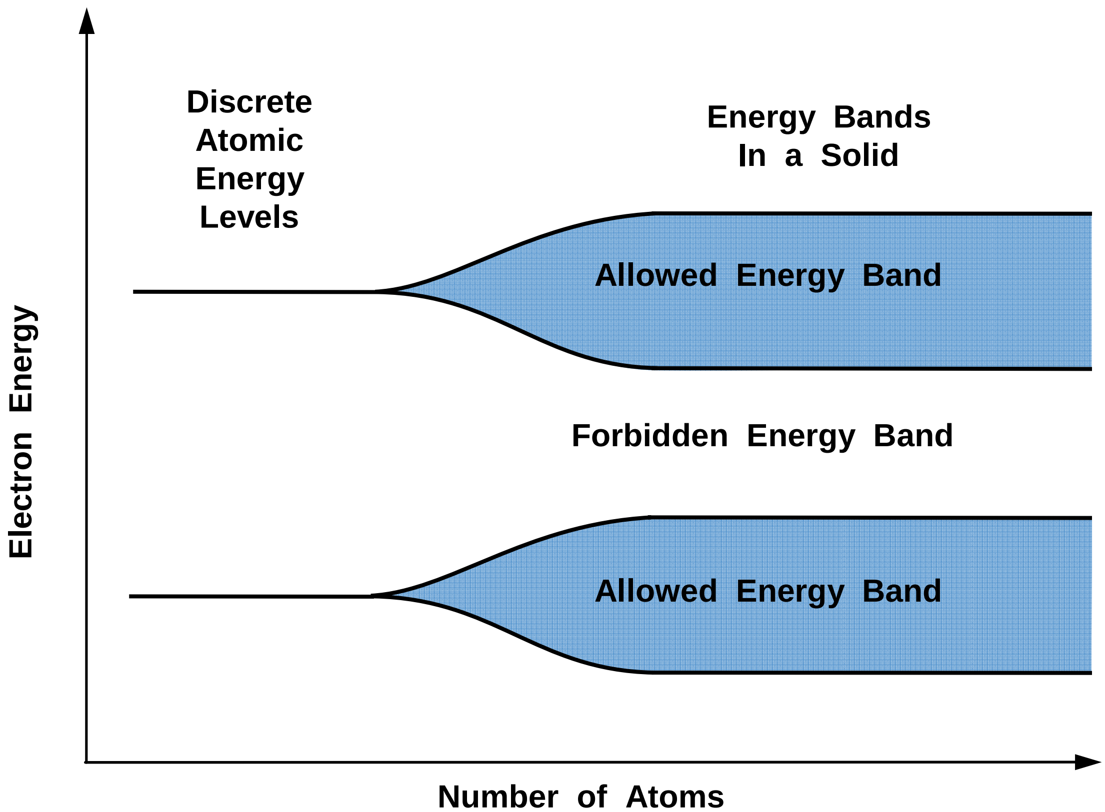

Figure 1-6: Electron Energy Levels In Discrete Atoms and Solids. (40280)

Figure 1-7 shows the electron energy levels in metals, insulators and semiconductors at a temperature of 0 Kelvin (-273.15C). In insulators, the lower energy band is filled and there is a large energy gap to the next higher (empty) energy band. There can be no net current flow in filled energy bands because there are no available states for electrons to move into. So insulators are insulating because of the large energy gap to the next available states and the electrons do not have enough energy to reach these states. In metals, the allowed bands overlap so electrons easily reach the higher states. In semiconductors there is a smaller energy gap and a small fraction of the electrons can reach the higher energy band at non-zero temperatures. The energy at which a state has a probability of being filled equal to 1/2 is referred to as the Fermi energy or Fermi level. For semiconductors without impurities, the Fermi level is in the middle of the energy gap. The filled lower energy band is called the valence band and the higher empty band is the conduction band. The energy difference between the conduction band minimum and the valence band maximum is referred to as the bandgap energy (E_g). When an electron moves from the valence band to the conduction band it leaves an empty state in the valence band. This empty state is treated as a positively-charged particle known as a hole. Holes respond to electric fields just like electrons except holes have a positive charge and a larger effective mass which is material dependent and approximately three times the electron effective mass.

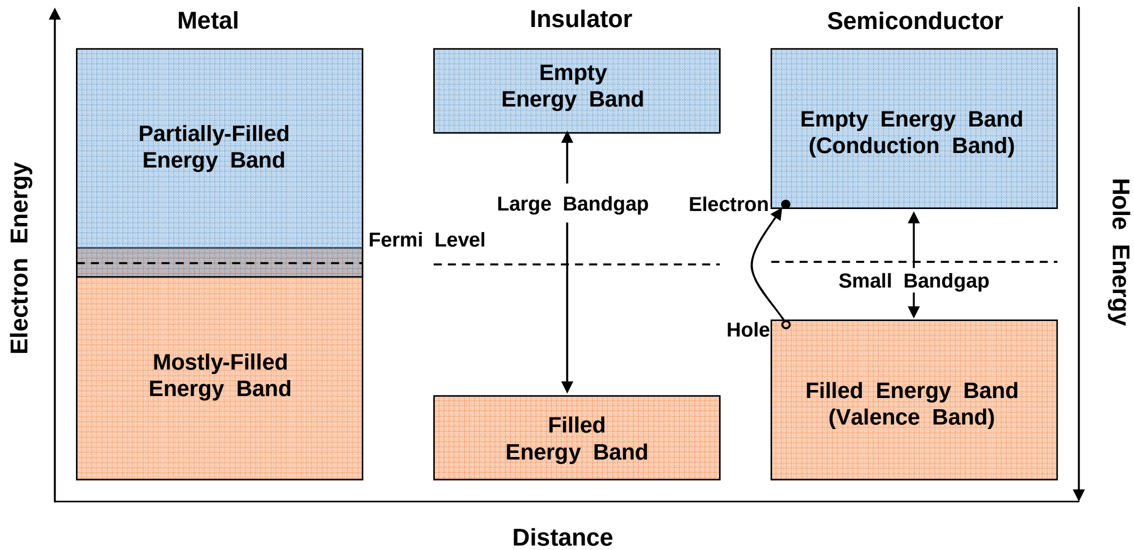

Figure 1-7: Electron Energy Levels In Metals, Insulators and Semiconductors At 0 Kelvin. (40281)

The Fermi level can be moved up or down by adding impurities, referred to as dopants, to the semiconductor. To make an n-type semiconductor, dopants with an extra valence electron are added. These added impurities add electrons to the conduction band and move the Fermi level closer to the conduction band. A p-type semiconductor is made by adding dopants with fewer valence electrons. These added impurities add holes to the valence band and move the Fermi level closer to the valence band. On the electron energy band diagrams, electrons tend to move to lower energy (flow down hill) by giving energy to the semiconductor. The hole energy is the negative of electron energy so on these diagrams, the hole energy decreases going up (increasing electron energy). Therefore, holes tend to move to higher electron energy (flow up hill) on these diagrams. These electrons and holes have a range of energies with the most carriers at the band edge and the number of carriers exponentially decreasing with increasing electron or hole energy.

1.3.3 Diodes and Light Emitting Diodes (LEDs)

A diode is made by putting an n-type semiconductor in contact with a p-type semiconductor forming a p-n junction as shown in the top part of Figure 1-8. In thermal equilibrium, the Fermi level must be constant across the junction so the conduction and valence bands need to bend in the transition from the n-type to the p-type semiconductor as shown in the middle part of Figure 1-8. Since electrons flow downhill, the electrons in the conduction band remain in the n-type half of the diode. In equilibrium, they do not have enough energy to move to the p-type side of the diode. Holes on the other hand flow uphill, so they are on the p-type side and do not have enough energy to move to the n-type side. In order for an electron in the conduction band to move to the valence band to recombine with a hole to emit a photon, they must overlap in space (horizontal axis in Figure 1-8). Forward biasing the diode (bottom of Figure 1-8) causes the electron region in the conduction band to extend further to the left and for the hole region in the valence band to extend further to the right. The two regions then overlap in the transition region which allows electrons and holes to recombine. This is referred to as spontaneous emission and devices that exploit this process are referred to as light emitting diodes (LEDs). The wavelength of the emitted light depends on the energy difference between the electrons and holes that recombine and is given by λ (µm) = 1.24 / E (eV), as discussed in Section 1.3.1. The wavelength is generally quoted as that in vacuum even though the light is in the semiconductor. To convert to the wavelength in the semiconductor, divide the wavelength in vacuum by the refractive index. The minimum energy difference is the semiconductor bandgap (E_g) giving the longest wavelength. Larger energy differences (shorter wavelengths) are also possible due to electrons and holes that are not at the band edges, as discussed above. Typical communications LEDs have a spectral width in the range of 30 nm to 150 nm [1-1].

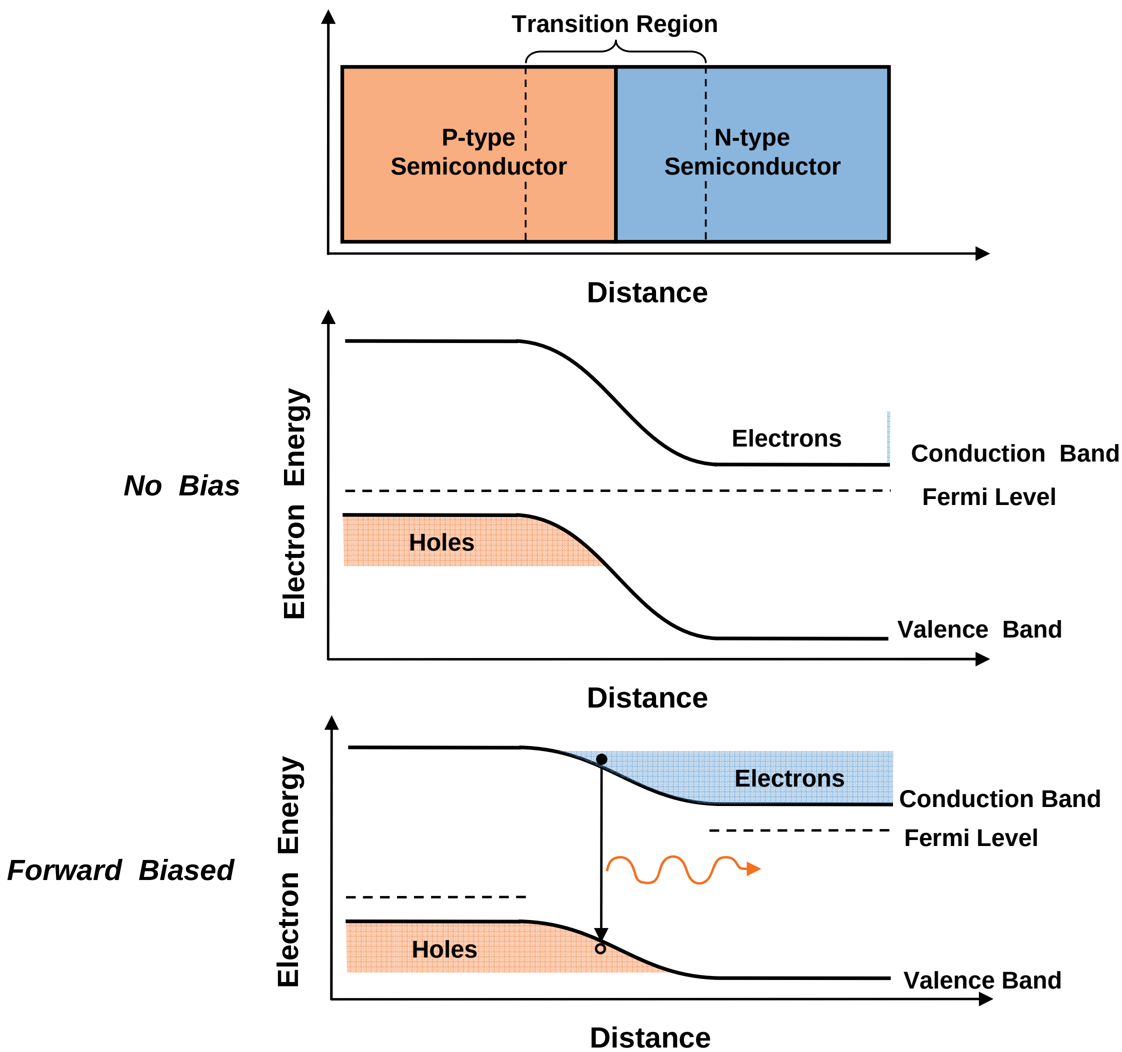

Figure 1-8: Band Diagram of An LED With No Bias and Forward Biased. (40282)

1.3.4 III-V Versus Silicon

In addition to conserving energy, the emission process must also conserve momentum. The photon momentum is approximately 3-4 orders of magnitude smaller than the electron momentum. Therefore, the electron's momentum essentially does not change when a photon is emitted and the electron moves vertically on the energy-versus-momentum diagram. Figure 1-9 shows simplified energy-versus-momentum diagrams for silicon and a typical III-V semiconductor. The III-V semiconductors (like GaAs and InP) have direct bandgaps which means that the conduction band minimum energy is directly above the valence band maximum in the energy versus momentum diagram. Therefore, the electrons can recombine with holes by emitting photons and still conserve energy and momentum. Silicon, on the other hand, has an indirect band gap which means that the conduction band minimum is not directly above the valence band maximum. So electrons cannot conserve momentum when recombining with holes by emitting photons. Momentum is conserved when the electrons give (or receive) momentum to (from) the semiconductor atoms (phonon emission or absorption) and then recombine with holes by emitting photons. The probability of this process occurring is very small. Therefore, in contrast to the III-V materials, silicon generally does not emit enough light to make a useful device and is not used for LEDs or lasers.

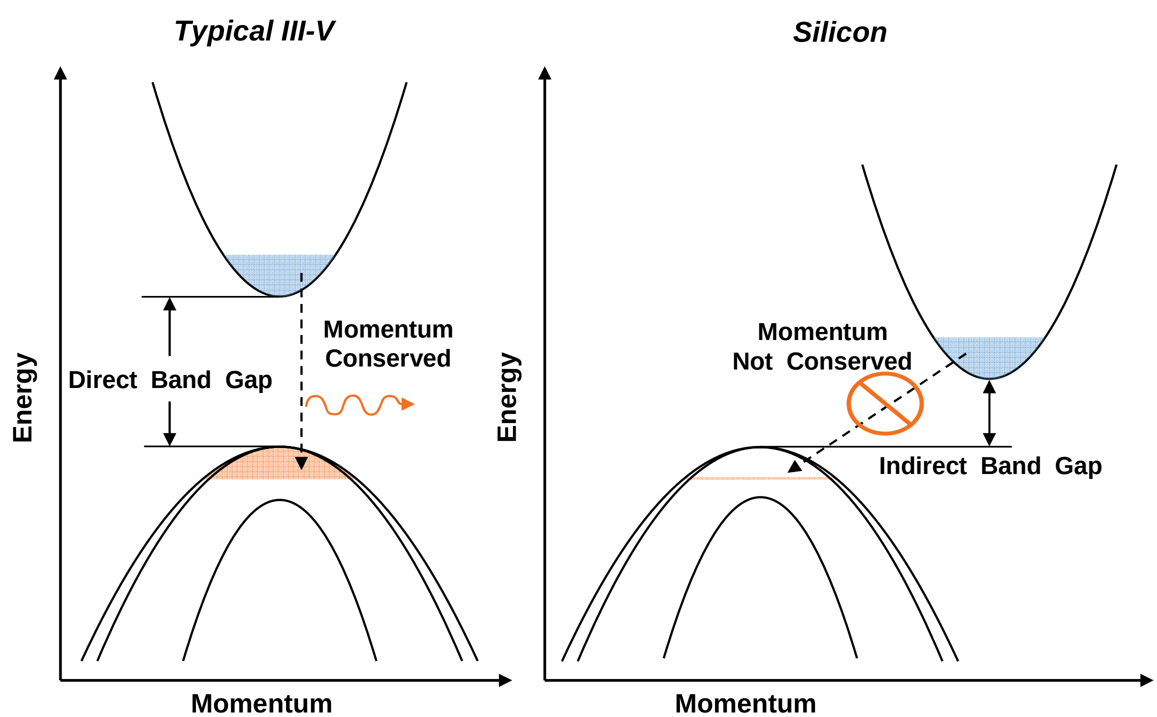

Figure 1-9: Energy Versus Momentum Diagrams for A Typical III-V Semiconductor and Silicon. (40283)

1.3.5 Lasers

A laser is made by starting with an LED and adding two parallel mirrors to it; one on each side of the LED. The distance between the mirrors is an integer number of half wavelengths of the light to form a resonant cavity. When the LED is turned on, some of the light from the spontaneous emission will reflect off of a mirror and pass through the LED again. When a photon passes near an electron in the conduction band, it can cause stimulated emission in which the electron recombines with a hole in the valence band and emits a photon that is coherent with, and traveling in the same direction as, the original photon. These photons are then reflected by the other mirror and pass through the material again causing more stimulated emission. This is referred to as optical gain. One of the mirrors is designed to be not perfectly reflecting and thus will allow some of the light to pass through; this pass through light forms the beam from the laser.

There are a few different techniques for making the mirrors in semiconductor lasers. The first is to cleave the semiconductor. Because of the crystalline nature of semiconductors, they cleave smoothly along crystallographic planes. There are two sets of crystallographic planes in a semiconductor wafer which are orthogonal to each other. The planes within each set are parallel to each other so two cleaves will form two parallel mirrors. The planes are also vertical so the light emission will reflect from one mirror to the other and stay within the semiconductor. These are referred to as Fabry-Perot lasers because the two mirrors form a Fabry-Perot cavity and they can also be referred to as edge-emitting lasers. The other main technique for making mirrors is to use a periodic variation of the RI (Bragg grating). Light is reflected at each RI change and when the period of the variation is a multiple of the wavelength of the light, constructive interference occurs. This technique can be used for edge-emitting lasers (distributed feedback [DFB] lasers and distributed Bragg reflector [DBR] lasers) or for vertical-cavity surface-emitting lasers (VCSELs). These will be discussed in more detail in Section 3.

1.3.6 Quantum Wells

In the discussion so far, we have talked about devices composed of a single semiconductor material, with separate regions of n-type and p-type doping. It is also possible to form a junction between two dissimilar semiconductors, such as GaAs and AlGaAs. In such a heterojunction, the Fermi level is constant across the junction but the conduction bands and valence bands will be offset because of the different band gaps of the two materials as shown in the left side of Figure 1-10.

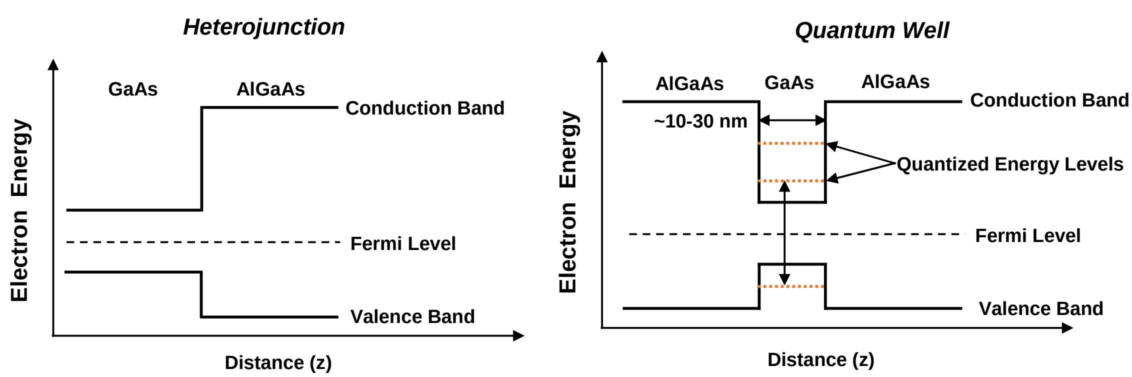

Figure 1-10: Band Diagrams for A Typical Heterojunction and Quantum Well. (40284)

A quantum well or double heterostructure results when a thin layer of narrow gap semiconductor is sandwiched between two layers of wider gap semiconductor, as shown in the right side of Figure 1-10. Any free electrons in these structures will be in the narrower gap material because its conduction band is at a lower energy. The structure becomes a quantum well when the narrow gap layer is thin enough (< ~50 nm) that the electron energy levels become quantized in this z-direction but are still continuous in the x- and y-directions. So rather than having a continuum of energy levels in the z-direction, the electrons and holes now have discrete allowed energy levels. For light emission in quantum wells, rather than moving from the conduction band edge to the valence band edge, electrons move from a quantized electron energy level to a quantized hole energy level. The energy difference is now larger than the band gap and increases as the layer is made thinner. Many of today's laser diodes utilize quantum wells in their structures to confine photon emission in the vertical direction. This will be discussed more in Section 3.

1.3.7 Coupling Light to Waveguides

In optical communication systems, the optical devices, such as lasers, detectors and different devices in planar lightwave technologies, need to be coupled to optical fibers. The coupling method used depends on economics, the beam divergence of the device and the spot sizes of the device and the fiber. For beams with a small divergence, such as VCSELs, it is possible to directly butt-couple the device to the fiber. For beams with a larger divergence, such as Fabry-Perot lasers, a lens system may be required. The lens could be as simple as a microsphere lens or a rounded or tapered fiber end or it could be a more complicated compound lens consisting of a pair of cylindrical lenses and a spherical lens. Drawings of these configurations are shown in Figure 1-11.

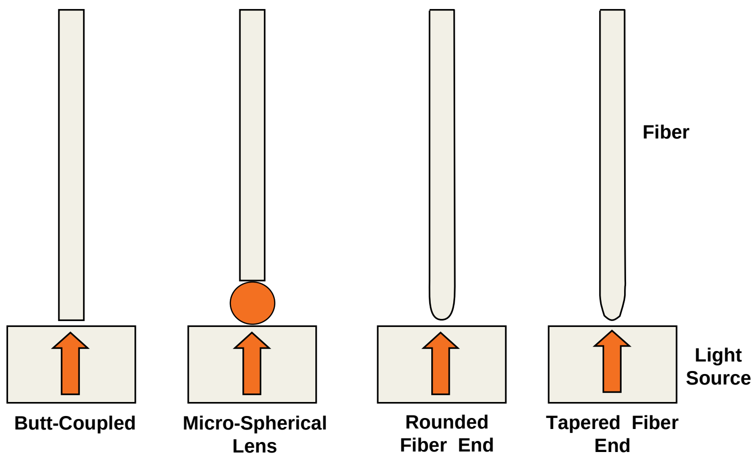

Figure 1-11: Fiber to Device Coupling Configurations. (40290)

1.4 Imparting Information on Light

In the previous section we discussed how to generate light for our system. Now we discuss how to modulate the light to encode data onto the light stream so that it will carry information. Here we will discuss the modulation schemes and in Section 3.3 we will discuss the hardware necessary to accomplish this.

1.4.1 Modulation Schemes

Figure 1-12 shows several common modulation schemes. The simplest modulation schemes (top 3 in Figure 1-12) use on-off keying (OOK) in which the light is either full on or is full off resulting in a series of light pulses in binary form. Non-return to zero (NRZ) is the most commonly seen approach in digital logic systems and can be directly used as the modulation signal. Return to zero (RZ) and non-return to zero inverted (NRZI) are different on-off based modulation approaches. Table 1-1 lists the encodings for these three on-off approaches and for Duobinary coding. Duobinary is a three-level modulation scheme that introduces controlled intersymbol interference (ISI) which results in a reduction of the bandwidth occupancy of the signal. There are several different implementations of duobinary. In the simplest implementation, the output level is determined by adding the levels of the current bit and the previous bit. For the example shown in Figure 1-12, with the first bit = ‘0’, the first output level is also ‘0’. The second output level is ‘1’ resulting from a ‘0’ in the first bit and a ‘1’ in the second bit. In this example, the first time the output level is a ‘0’ is at the sixth bit when there are two consecutive ‘0’s (5th and 6th bits) and the first time the output level is a ‘2’ is at the eighth bit when there are two consecutive ‘1’s (7th and 8th bits). When this scheme is implemented optically, a ‘1’ level is transmitted as no light and ‘0’ and ‘2’ levels are 100% light transmission with opposite optical phases.

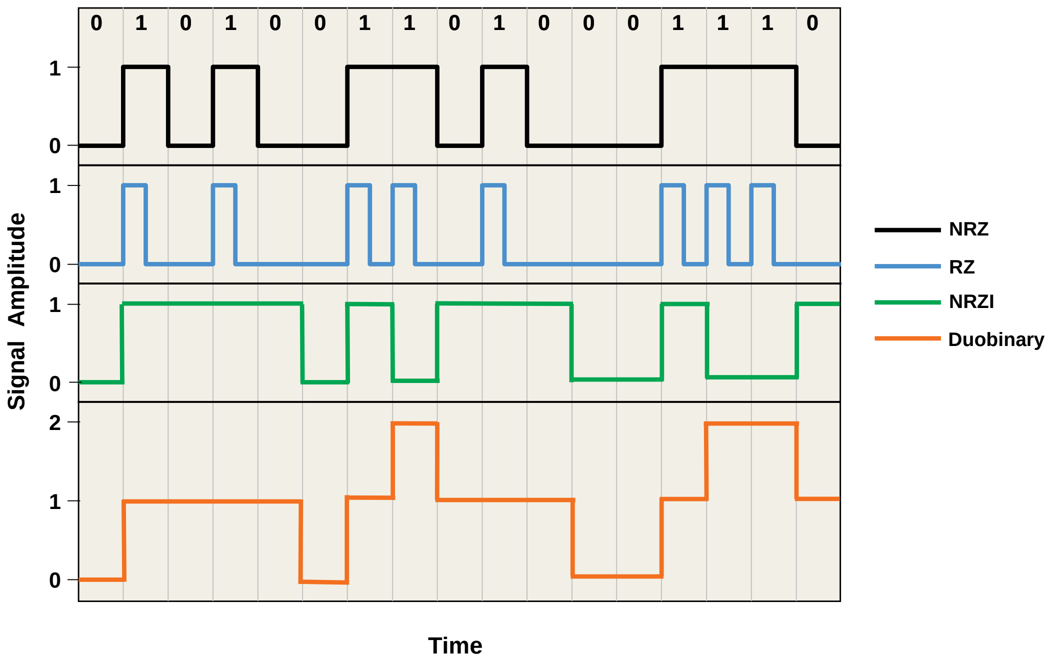

Figure 1-12: Bit Streams of Several Common Modulation Schemes. (40291)

Modulation Scheme	**Rules**
Non-Return to Zero (NRZ)	0=off, 1=on
Return to Zero (RZ)	0 = off, 1 = on for half of bit duration then off
Non-Return to Zero Inverted (NRZI)	0 = no change in amplitude 1 = transition from on to off or off to on
Duobinary	Level = previous bit + current bit 0 = 0 + 0 1 = 0 + 1 or 1 + 0 2 = 1 + 1

Table 1-1: Rules for Several Common Modulation Schemes.

Other modulation schemes are also being used. One is frequency shift keying (FSK) in which information is encoded using discrete frequency changes. The top portion of Figure 1-13 shows the simplest implementation which is binary FSK where a 1 is represented by one frequency and a 0 by another frequency. Another modulation scheme is phase shift keying in which information is encoded by changing the phase of a carrier signal. Again the simplest form is binary phase shift keying (BPSK) which uses two phases so the phase difference between a 1 and a 0 is 180°. Figure 1-13 shows example signals for these two modulation schemes. In this figure, only two to three cycles are shown for each bit. In reality, at approximately 1550 nm, there will be tens of thousands of cycles per bit and phase transitions will occur over thousands of cycles. For convenience and to make the illustrations clearer, only a few cycles will be shown in the figures for all of the phase shift keying schemes in this section.

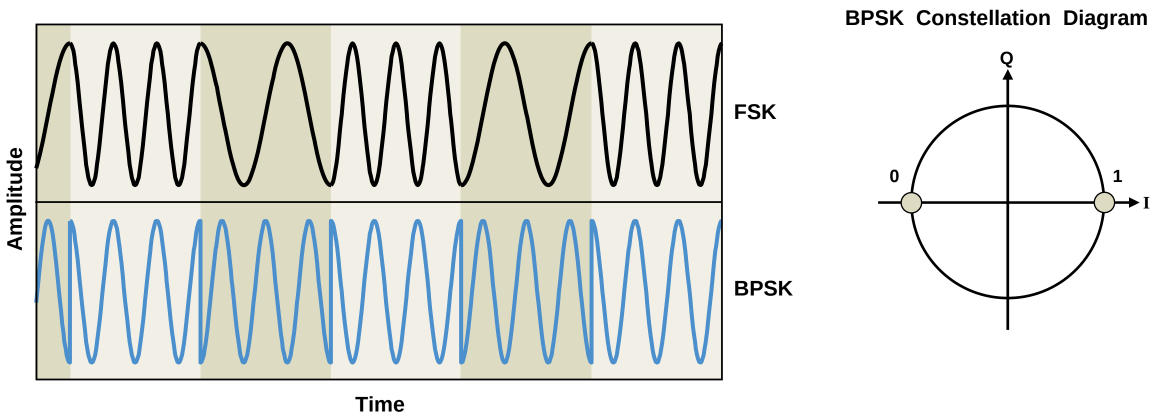

Figure 1-13: Waveforms for FSK and BPSK. (40292)

A convenient way to represent PSK schemes is on a constellation diagram where the x- and y-axes are termed the in-phase and quadrature axes respectively due to their 90° separation. The amplitude of each point along the in-phase axis is used to modulate a cosine (or sine) wave and the amplitude along the quadrature axis to modulate a sine (or cosine) wave. The constellation points are generally positioned with uniform angular spacing around a circle to give maximum phase separation between points. Figure 1-13 also shows the constellation diagram for BPSK.

Quadrature phase shift keying (QPSK) uses four phases to encode the information and therefore transmits 2 bits per symbol. This scheme can be viewed as two independently modulated quadrature (90° out of phase) carriers. With this interpretation, the even (or odd) bits are used to modulate the in-phase component of the carrier, while the odd (or even) bits are used to modulate the quadrature-phase component of the carrier [1-2]. Figure 1-14 shows the in-phase and quadrature components along with the overall signal. In this example, the odd bits modulate the in-phase component (I) and the even bits modulate the quadrature component (Q). The constellation diagram for QPSK is shown on the right side of this figure.

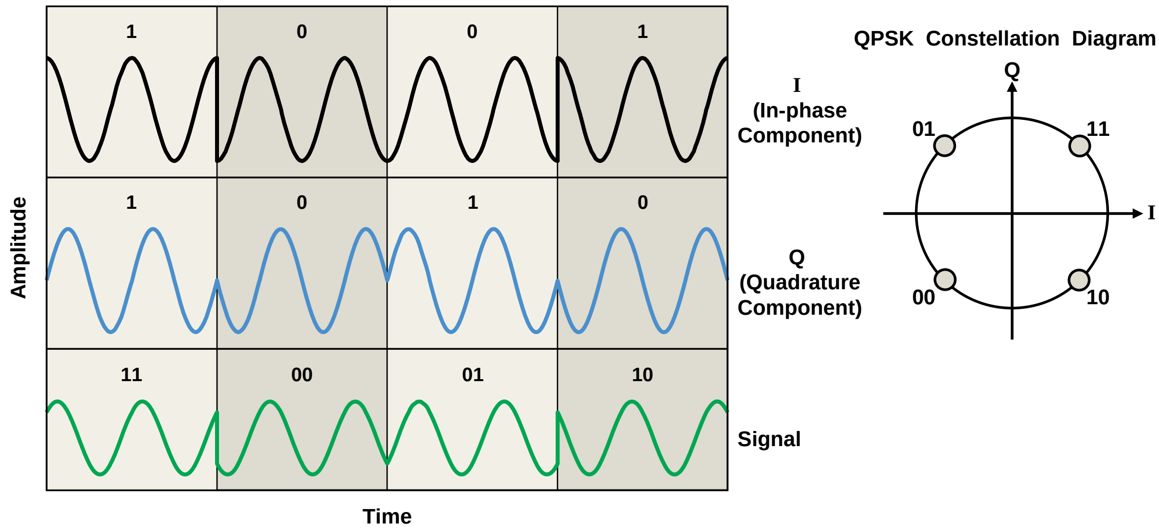

Figure 1-14: Waveforms for QPSK. (40293)

BPSK and QPSK are both subject to errors if the constellation is rotated, i.e. if there is a phase shift, during transmission. To overcome this problem, the encoding is done such that the data is used to change the phase rather than set the phase. This is referred to as differential BPSK (DBPSK) and differential QPSK (DQPSK). Example bit streams for these two modulation schemes are shown in Figure 1-15. In DBPSK, a ‘1’ can be represented as a 180° phase change from the previous bit and a ’0’ as a 0° phase change. In the example in Figure 1-15, the first two bits are ‘1’s so there are 180° phase changes for these two bits. The next three bits are ‘0’s, so there are no phase changes for these. The next two bits are ‘1’s so there is a 180° phase change for each of these bits. Finally, the last bit is a ‘0’ and there is no phase change. In DQPSK, the phase changes are 180°, 0°, +90° and -90° for bits 11, 00, 01 and 10, respectively. In the example in Figure 1-15, the first two bits are 11, so there is a 180° phase change from the previous bits. The next bits are 00, so there is no phase change. The last two pairs of bits are 01 and 10. For these, the phase changes are +90° and -90°, respectively.

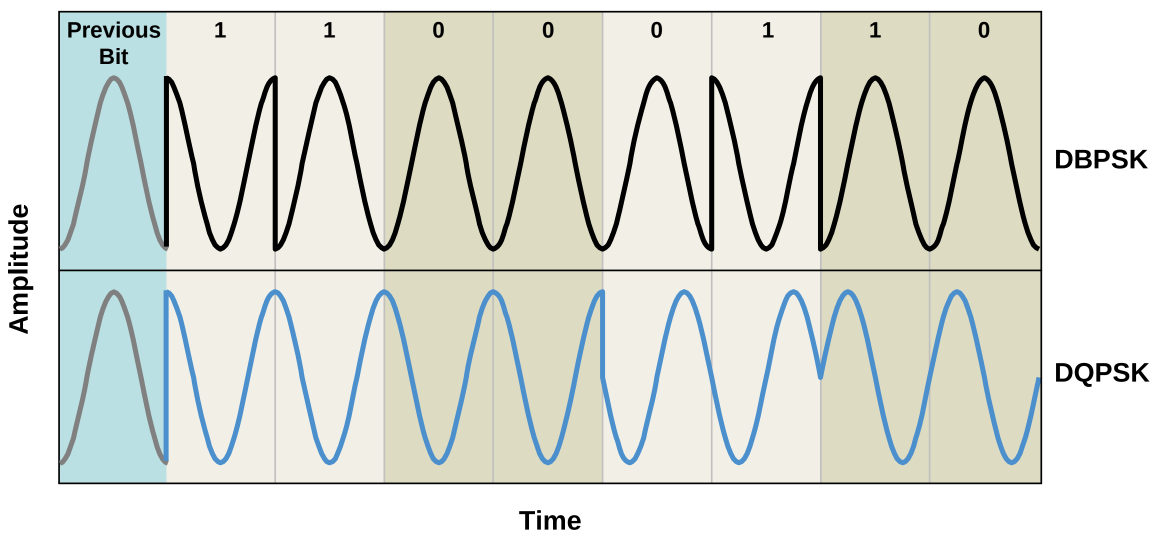

Figure 1-15: Waveforms for DBPSK and DQPSK. (40294)

1.4.2 Extinction Ratio and Optical Modulation Amplitude

Next we will discuss several characteristics of modulators. Two important specifications for laser transmitters related to modulation are the extinction ratio (ER or r_e) and the optical modulation amplitude (OMA). Figure 1-16 shows a typical eye diagram showing the power levels for '1' and '0' bits. P1 denotes the power level for a '1' and P0 denotes the power level for a '0'. The extinction ratio is defined as the ratio between the powers of the '1' level and the '0' level

$$ER = P1 / P0$$

and OMA is defined as the difference between the powers of the '1' level and the '0' level

$$OMA = P1 – P0.$$

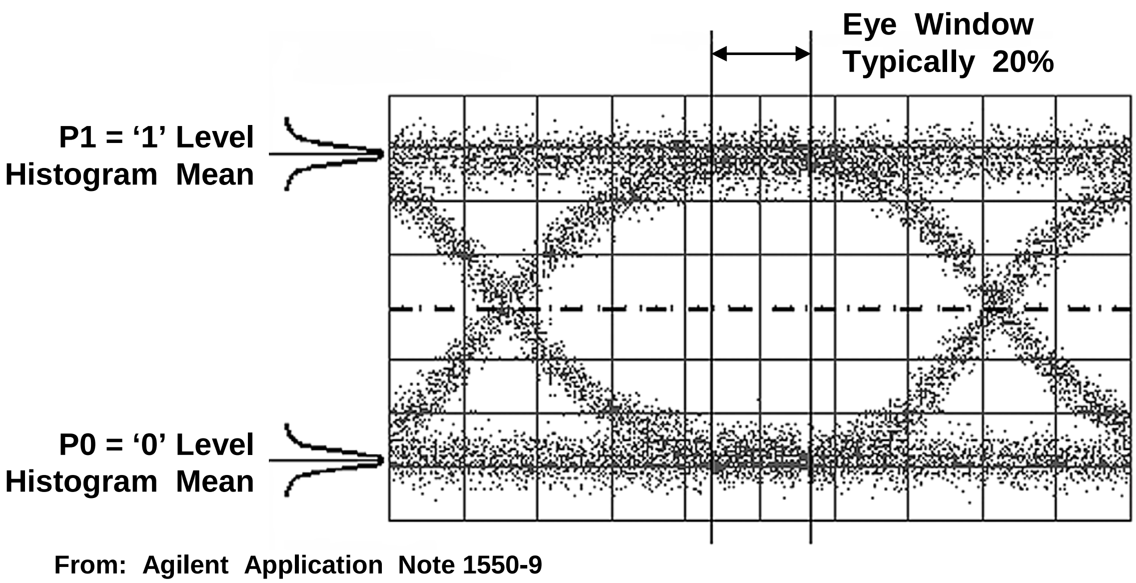

Figure 1-16: Typical Eye Diagram At the Modulator Output. (40295)

Ideally P0 is zero and ER is infinite but in reality P0 is greater than zero. The bit error rate (BER) of a link is a function of ER; as ER increases, the BER of a link decreases. So when the extinction ratio is not infinite, the transmitted power must be increased in order to maintain the same bit error rate (BER). This increase in transmitted power due to non-ideal values of extinction ratio is called the "power penalty" [1-3]. The power penalty for a given ER is shown in Figure 1-17.

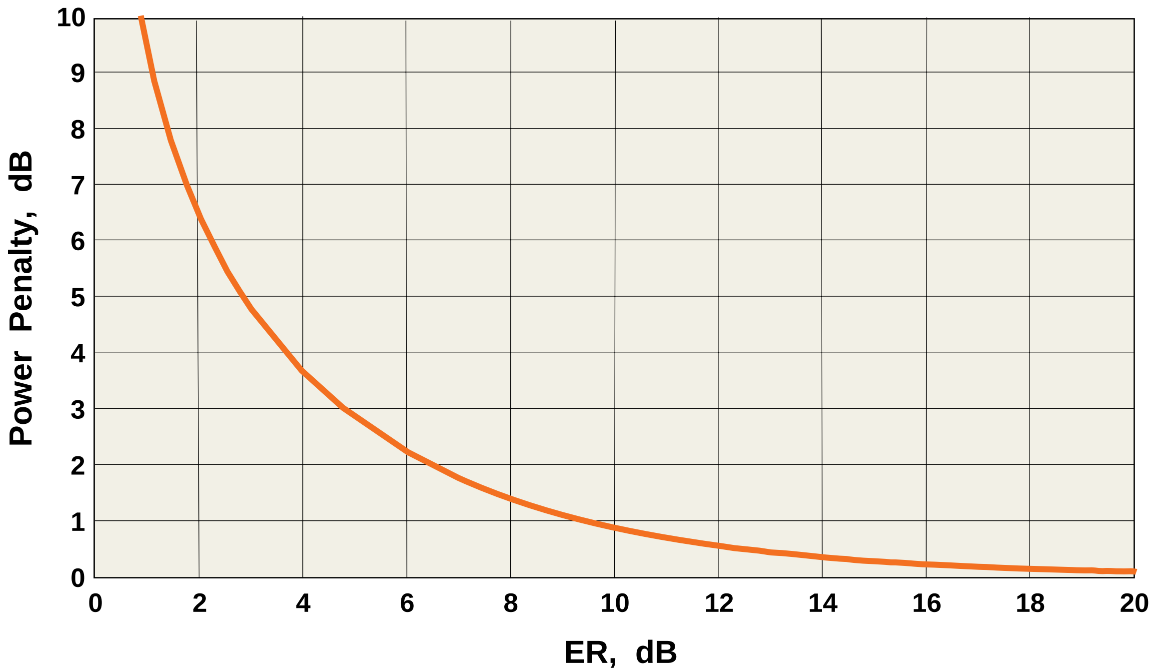

Figure 1-17: Power Penalty As A Function of Extinction Ratio. (40296)

1.4.3 Modulation Approaches

The light in simple on-off keying modulation schemes is generally modulated in one of two ways. The first is direct modulation in which the laser current is modulated by the OOK signal to generate a modulated light stream. VCSELs and some DFB lasers are examples of directly-modulated lasers (DMLs). The second approach is to use a separate modulator to modulate the continuous wave (CW) laser. In general any laser can be used with an external modulator and thus become an indirectly modulated laser. One example is the electro-absorption modulated laser (EML) which is a DFB laser integrated with an electro-absorption modulator (EAM).

Both types of modulation have their advantages and disadvantages. Directly-modulated lasers are less complicated because a separate modulator is not required, but they suffer from dynamic effects on the emitted spectrum, such as changes in peak wavelength, spectral bandwidth and the amplitude of the individual cavity modes. Also for higher speed operation, the lasers are not turned all the way off for the ‘0’ bit level resulting in lower ER and OMA. The indirectly-modulated lasers are more complicated due to the modulator but they operate at higher data rates (modulation above 100 Gbps is possible) and one can employ a CW laser with low relative intensity noise (RIN), such as a DFB laser, so that there is less noise in the light output. Different types of modulators are discussed in more detail in Section 3.3.

1.5 Light Detection

1.5.1 Role of Photodetectors

The role of the photodetector is to convert an optical signal to an electrical signal. The three basic processes are 1) carrier generation by incident light, 2) carrier transport to electrical contacts, and 3) interaction of current with the external circuit to provide the output signal [1-4].

A fourth process that is present in some devices is carrier multiplication by whatever current-gain mechanism may be present.

1.5.2 Photon Absorption

Carrier generation by incident light is essentially the reverse process of light emission in semiconductors. An incident photon with energy greater than the semiconductor bandgap energy is absorbed by moving an electron from the valence band to the conduction band and leaving behind a hole in the valence band as shown in Figure 1-18. Light with a wavelength shorter than the cutoff wavelength (= $hc/E_g = 1.24/E_g$) will be absorbed while longer wavelength light will pass through the semiconductor. Despite silicon's indirect bandgap which prevents light emission, it can be used as a photodetector. To conserve momentum, the electron being excited from the valence band to the conduction band interacts with the silicon atoms to gain the momentum needed to reach the conduction band minimum. The probability of this process occurring is high enough to make a useful device but at reduced efficiency compared to direct bandgap semiconductors. Figure 1-19 shows the absorption coefficients for several materials commonly used to make photodetectors. The absorption coefficient (α) is used to calculate the penetration depth of light in the material using the Lambert-Beer law:

$$I = I_0 \exp(-\alpha x),$$

where I is the intensity in the material, I_0 is the incident intensity and x is the depth. A higher absorption coefficient corresponds to higher absorption near the surface of the material. As can be seen in the figure, the absorption coefficient of GaAs is approximately 16 times that of Si at a wavelength of 850 nm.

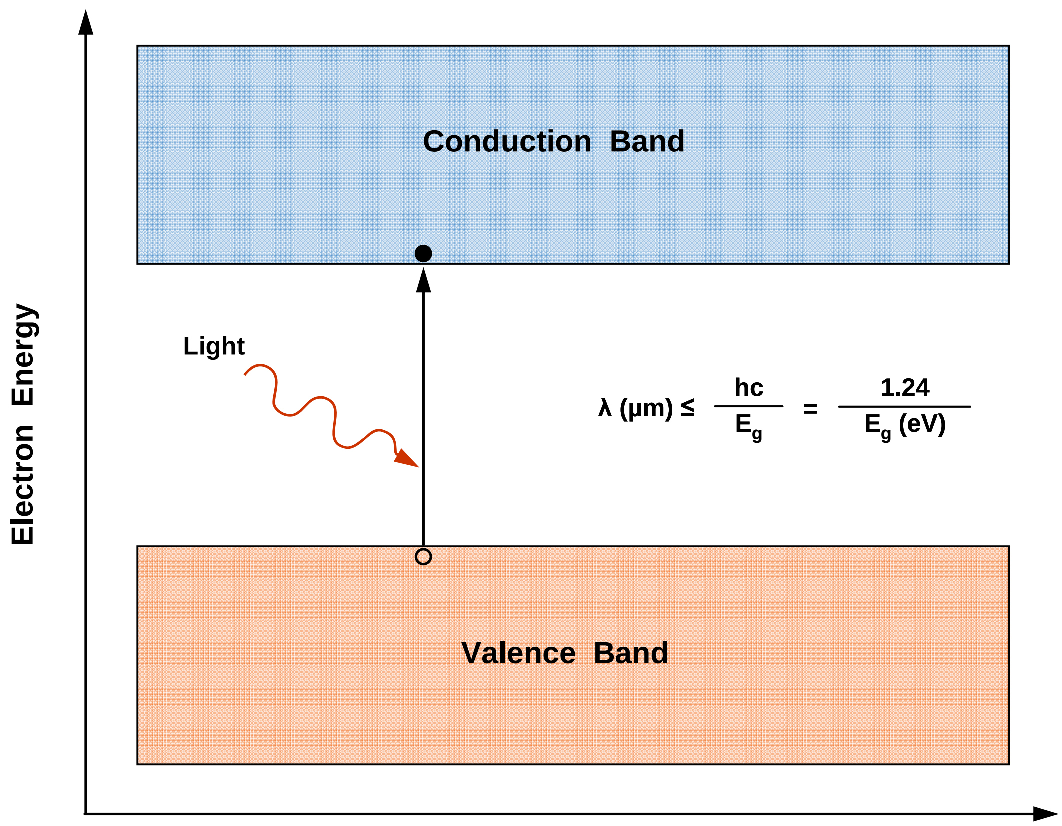

Figure 1-18: Illustration Showing Light Absorption In A Semiconductor. (40310)

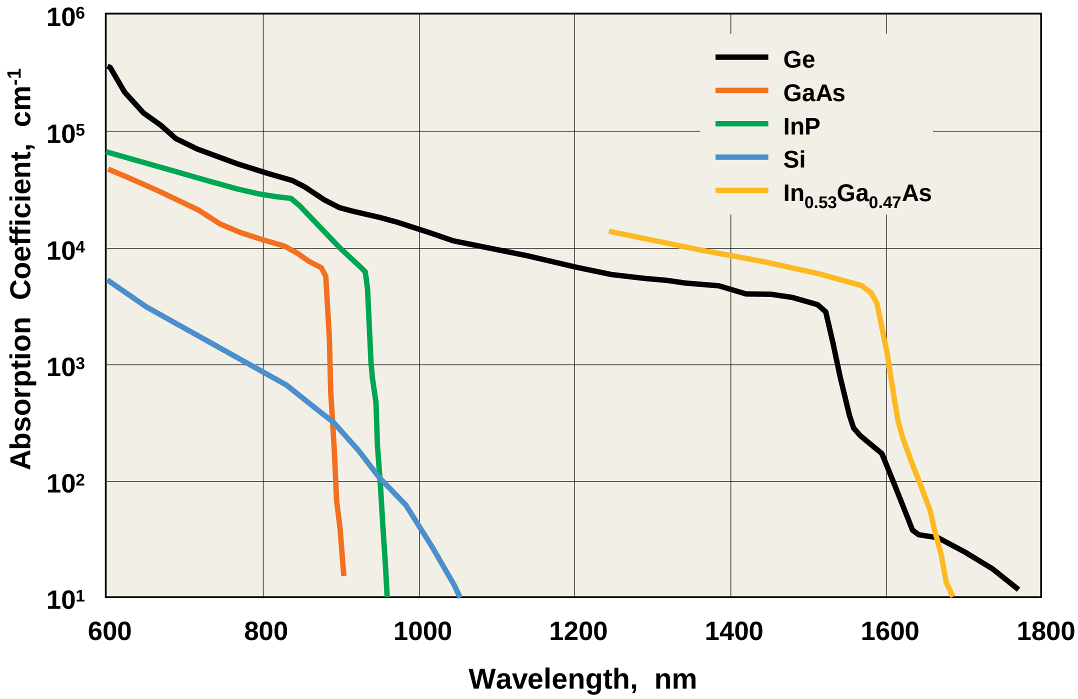

Figure 1-19: Absorption Coefficients of Several Photodetector Materials. (40335)

1.5.3 Photodiodes

The photo-generated carriers discussed in the previous subsection are transported to the contacts using the electric field in the depletion region of a reverse biased diode. Photodiodes are generally used because the electric field generated at the P-N junction in a reverse biased diode effectively moves the electrons and holes to the electrical contacts. Several photodiode characteristics are:

Quantum efficiency (η) = number of electro-hole pairs generated per photon,

Responsivity (r) [A/W] = ratio of the output current to the incident optical power (formerly referred to as sensitivity).

Electrons in the valence band can also gain enough thermal energy to move into the conduction band, resulting in an electron-hole pair without the presence of light. These charge carriers are then swept to the contacts by the electric field resulting in what is known as dark current. The dark current has a strong dependence on the temperature and also has an inverse dependence on the bandgap of the semiconductor. It is a performance limiting characteristic for the photodetector since small photon fluxes will be masked by the dark current. Variations on the photodiode, which will be discussed in the following paragraphs, are the PIN diode and the avalanche photodiode (APD).

1.5.4 PIN Diodes

The P-N junction depletion region of a diode has a relatively small volume for absorbing photons so an undoped or lightly doped layer, referred to as an intrinsic layer, is added between the P-type and N-type layers. Hence the name P-I-N diode. The equilibrium and reverse biased band diagrams of this device are shown in Figure 1-20. The region with an electric field for separating photo-generated electrons is much larger and the probability of absorbing photons is increased. The P-type and N-type regions are made relatively small because photo-generated electrons and holes do not experience an electric field and are more likely to recombine with each other giving their energy to the semiconductor atoms as heat.

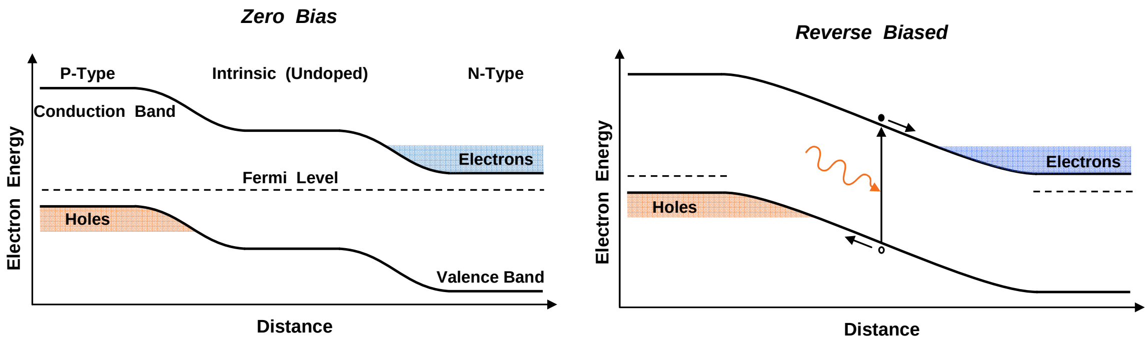

Figure 1-20: PIN Diode Band Diagram In Equilibrium and Reverse Biased. (40311)

1.5.5 Avalanche Photodiodes (APDs)

The APD is similar to the PIN diode except it is operated at large reverse biases (30-300 V) so that it is near the reverse breakdown voltage. The large bias results in a high electric field in the junction region. Relatively thin n+ and p layers form the junction and the intrinsic layer is lightly doped p-type and is referred as the π region as shown in Figure 1-21. Electron/hole pairs are photo-generated in the π-region and the electric field causes the electrons to drift toward the junction. Here they are accelerated in the high electric field and acquire enough energy to dislodge additional electron-hole pairs in collisions with the semiconductor atoms in a process known as impact ionization. These new electrons and holes are then accelerated and cause more impact ionization and eventually result in an avalanche multiplication effect (photoelectric current gain) where large numbers of electrons and holes are generated. This amplification can be as high as 10 – 100 times and gives APDs a higher responsivity than PIN diodes. The multiplication factor has a strong dependence on temperature and voltage so these need to be stabilized. Dark current is also multiplied in these devices resulting in higher noise levels than PIN photodiodes.

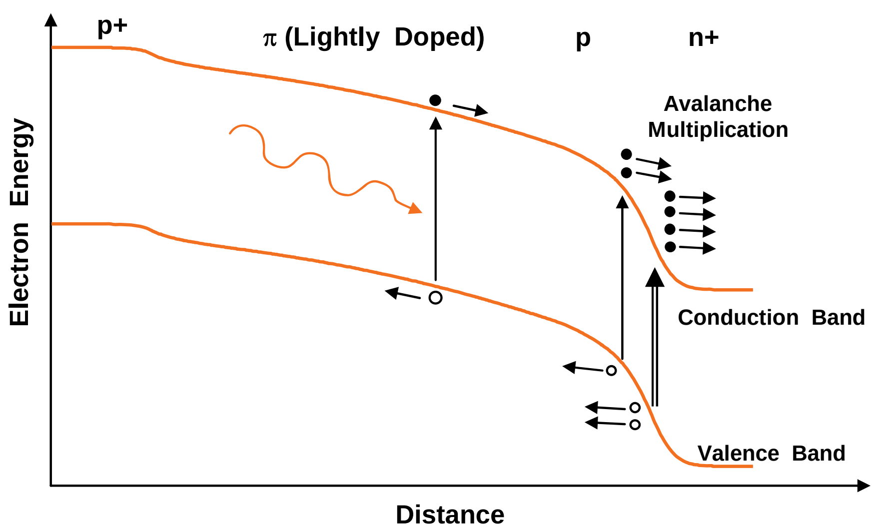

Figure 1-21: APD Band Diagram Under Reverse Bias. (40312)

1.5.6 Metal-Semiconductor-Metal (MSM) Photodiodes

Another type of photodiode is the Schottky-Barrier or Metal-Semiconductor-Metal (MSM) photodiode. In these photodiodes, one side of the P-N junction is replaced with a Schottky metal. The resulting band diagram is shown in Figure 1-22. Electron/hole pairs are photogenerated in the n-type semiconductor when photons are absorbed. The electrons reach the Schottky metal by thermionic emission over the barrier or tunneling through the barrier. These photodetectors have the potential to be faster than P-N junction based photodiodes but they have not found widespread use yet.

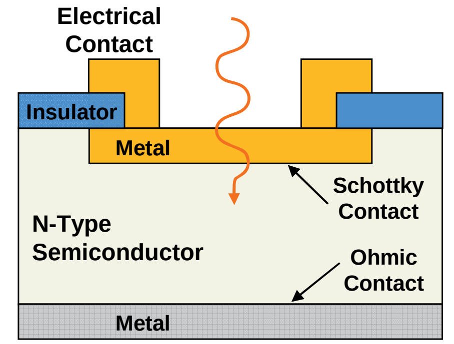

From: Understanding Optical Communications By H. J. R. Dutton

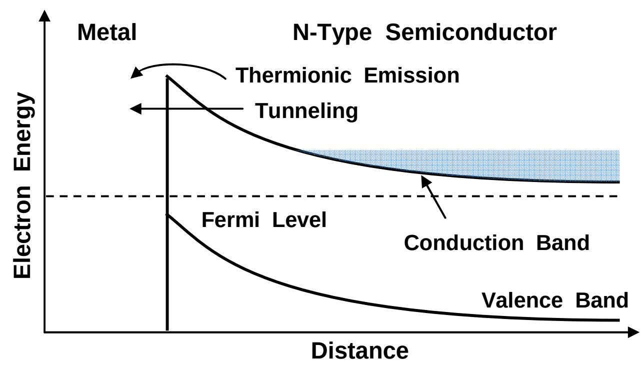

From: Low-Dimensional Semiconductors By M. J. Kelly

Figure 1-22: MSM Structure and Band Diagram. (40313)

1.5.7 Transimpedance Amplifiers (TIAs)

The output of the photodetector is basically a current source that needs to be converted to a voltage source before being amplified. This is accomplished using a transimpedance amplifier. Figure 1-23 shows the block diagram incorporating an op amp and a feedback resistor. The figure also shows the schematic diagram for a differential implementation using bipolar transistors. For systems operating at high speeds (>20 Gbps), TIAs are typically made using GaAs, InP or SiGe technologies.

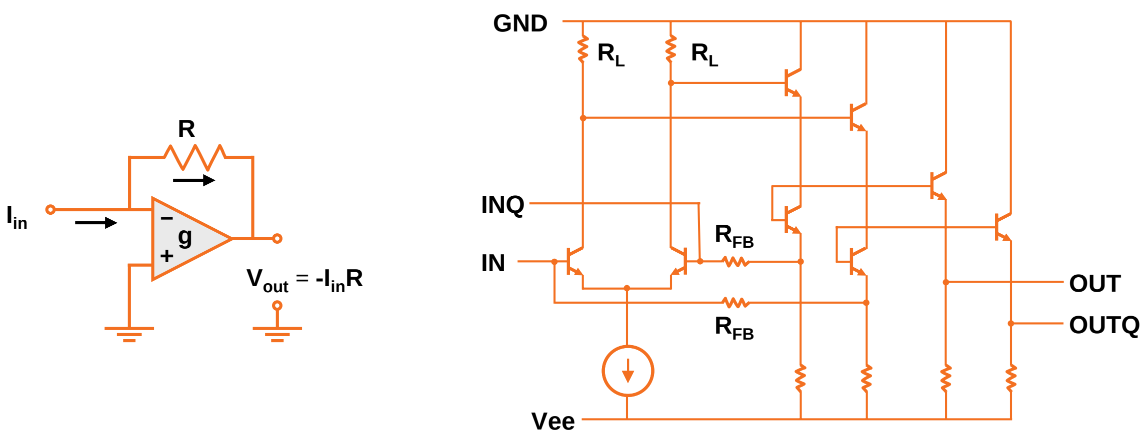

Figure 1-23: Block Diagram and Schematic Diagram of A TIA. (40314)

1.6 Multiplexing/Demultiplexing

One way to transmit more data over a link is to multiplex several signals onto the same link. Several of these techniques will be discussed in the following section.

1.6.1 Time Division Multiplexing (TDM)

One technique for multiplexing data is time division multiplexing (TDM) in which several lower speed bit streams are simply muxed together to form a higher speed data stream. Each channel is assigned a fixed-length time slot on the TDM data stream and its data is put on the combined data stream during its time slot. The time slots can be as short as one bit or as long as hundreds of bits. Figure 1-24 shows an example with four channels (with all data bits set to 1) at the same data rates TDM'ed onto a single bit stream. The individual channels do not have to be at the same data rates. In cases where they are not at the same rate, the higher rate channels will be assigned to more than one time slot. This technique can be used to combine electrical TDM (ETDM) or optical TDM (OTDM) signals. The advantages of OTDM are that lower speed and lower power electronic circuits can be used and data rates higher than those available in electronic circuits can be achieved.

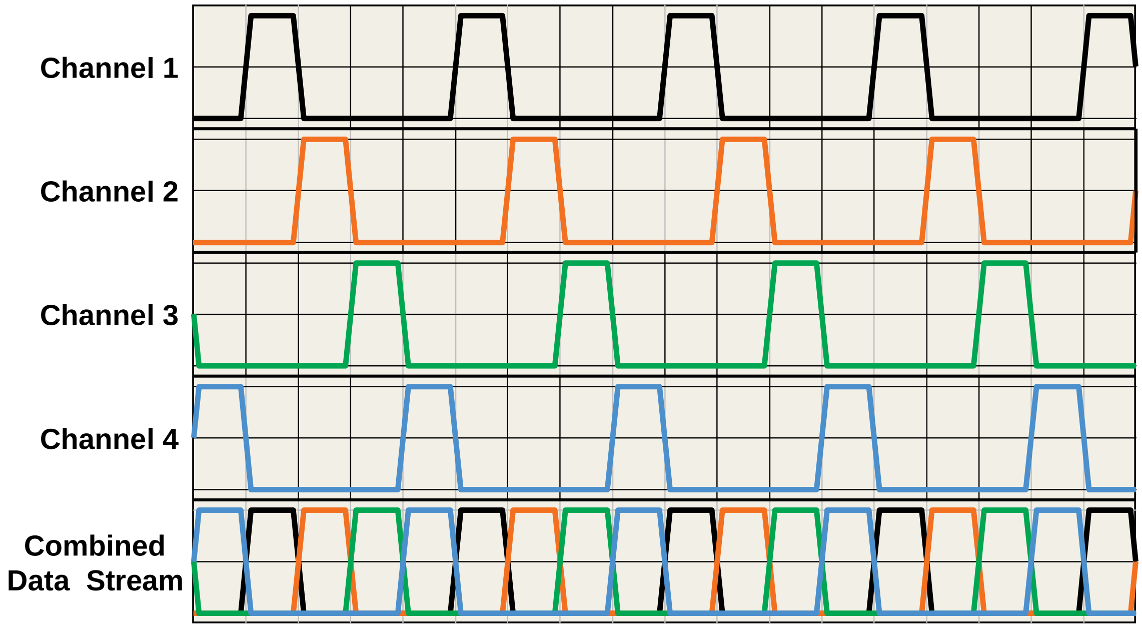

Figure 1-24: Idealized Input and Output Data Streams for 4-Channel TDM System. (40315)

1.6.2 Wavelength Division Multiplexing (WDM)

Another technique used to carry more information over a given fiber is wavelength division multiplexing (WDM); in which each multiplexed channel is given its own wavelength (color) of light on the fiber. The channels are independent so they can be at different data rates, use different encoding schemes and protocols and can even travel in opposite directions. The number of channels and channel spacing varies depending on system requirements. Some systems have been designed for coarse WDM (CWDM) with two channels; one at 1550 nm and the other at 1310 nm. Dense WDM (DWDM) systems have been demonstrated with up to 160 channels with channel spacing of 0.2 nm (25 GHz). There is not an agreed upon definition of ‘dense,’ so systems with channel spacing less than approximately 3.5 – 4 nm have been referred to as DWDM. The International Telecommunications Union (ITU) has defined 18 CWDM channels from 1270 nm to 1610 nm with 20 nm spacing (2,500 GHz) (ITU G.694.2) and 73 DWDM channels from 1519.5 nm to 1577 nm with an average spacing of 0.8 nm (100GHz) (ITU-T G.694.1). Generally DWDM systems use wavelengths in the 1540 – 1560 nm range to take advantage of the lower loss in optical fibers in this range and to use erbium doped fiber amplifiers (see Section 3.4.1).

In a WDM system, additional components are needed to multiplex the different channels onto a single fiber at the transmitters and to demultiplex the channels onto individual fibers at the receivers. Figure 1-1 shows a system using AWGs to perform the muxing and demuxing. These are discussed in detail in Section 1.6.3. Simple Y-couplers or other types of grating-based couplers could also be used, depending on the application.

One of the critical issues in DWDM systems is the wavelength stability of the lasers. The wavelength of the laser must stay within its allotted channel (wavelength range) to avoid cross-talk with adjacent channels. Generally, the lasers are temperature stabilized using a Peltier device (sometimes referred to as a thermoelectric cooler – TEC) to heat or cool the laser as needed. This adds to the overall power consumption of the link. For long-haul telecommunications applications, this is a relatively small amount of power and additional cost. However, for HPC applications this can be a significant portion of the power needed to transmit and receive a bit. Another source of wavelength variation is chirp, which results from RI variations in the laser cavity due to changing carrier flux densities when the laser is turned on. To avoid this, external modulators are used rather than directly modulated lasers.

There are several non-linear effects that can affect optical fiber transmission, especially at longer distances and at higher optical power levels, as typically seen in DWDM systems. “Longer distances” is again an ambiguous phrase and it is difficult to find any hard numbers in the literature. However, it is generally used in the context of long-haul telecommunications which implies distances well over one kilometer. Because of the much shorter distances and moderate power levels in HPC applications, it is most likely that these non-linear effects will not be significant. For completeness, however the main non-linear effects are described below.

Four-Wave Mixing (FWM) – This effect occurs when two or more wavelengths of light travel in the same direction and with the same phase on a fiber. The signals mix to produce new signals at wavelengths with the same spacing as the mixing signals. Dispersion causes a phase change between the two original wavelength signals and reduces FWM. So in these cases, a small amount of dispersion is good [1-1].

Stimulated Brillouin Scattering (SBS) – This effect occurs with narrow linewidth (~0.1 GHz) lasers at relatively high power levels (several mW). The electromagnetic field of the light wave causes mechanical vibrations in the fiber resulting in periodic RI changes in the fiber forming a moving diffraction grating. Then, over a long interaction length, the light is reflected by the diffraction grating. For HPC DWDM applications, the fiber is expected to be too short for this effect to be significant and the signal linewidth is expected to be too large (~25-100 GHz) [1-1].

Stimulated Raman Scattering (SRS) – This effect is similar to SBS, but it occurs when the light interacts with vibrating molecules. The light can be scattered in both the forward and backward directions and generally results in a transfer of power from a shorter wavelength signal to a longer wavelength signal. SRS increases with increased power, so keeping the channel power below approximately 3 mW reduces this effect [1-1].

1.6.3 Mux/Demux Components

In a WDM link, optical components are needed to take several channels with different wavelengths on separate fibers and combine them onto a single fiber (multiplex) and to take several channels with different wavelengths on a single fiber and split them onto multiple fibers based on their wavelengths (demultiplex). There are several different components available with different numbers of input and output channels and loss characteristics. They are typically based on either optical fibers or planar lightwave technology. In many cases these mux/demux components are referred to as couplers.

The first component to be discussed is a simple resonant coupler made in either fiber or planar technologies. The coupler consists of two fiber cores that are brought together for a specific coupling length. When two fiber cores are brought within approximately 5 µm of each other, all of the light entering one fiber will be transferred to the other fiber and then back to the first. This oscillation between fibers continues for the length of the coupled section. This is illustrated in Figure 1-25. The coupling length is the distance required for 100% of the light to transfer from one fiber to the other. The coupling length depends on the gap between the fiber cores and on the wavelength of the light. If the coupled section is half of the coupling length, half of the light will exit from each fiber. This is known as a 3 dB coupler. The insertion loss can be as low as 1.2 dB and the isolation (power from one input being reflected back to the other input) is often better than 40 dB. In planar technology, a 3-port y-coupler (one input waveguide and two output waveguides), with behavior similar to the 3 dB coupler, can be built.

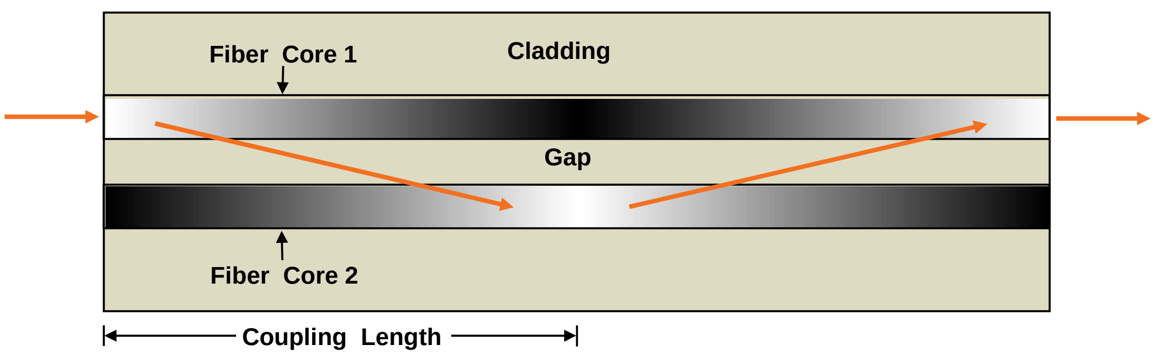

Figure 1-25: Diagram of the Operation of A Fiber Resonant Coupler. (40316)

1.6.4 Star Coupler

Couplers with more than two inputs and two outputs where each input signal is coupled to each output fiber are known as star couplers. There are two configurations for star couplers; inputs and outputs on separate fibers (transmissive star coupler), and inputs and outputs on the same bidirectional fiber (reflective star coupler). The output power in each fiber is much lower than the input power because it is evenly split between all of the output fibers. There is an additional 3 dB of loss for every doubling of the number of outputs. For a star coupler with 32 inputs and 32 outputs, the minimum loss is 15 dB. Star couplers have been fabricated with 2 dB of loss (excess loss) beyond the theoretical minimum loss.

One method of making a transmissive star coupler is similar for fiber and planar technologies, as shown on the left side of Figure 1-26. The input and output fibers or waveguides are attached to a central free space region which consists of a region of higher RI material surrounded by lower RI material. For fiber-based star couplers, this region consists of fused fibers and for waveguide-based star couplers, it consists of a slab of waveguide material (e.g. quartz, Si or InP). Light enters on one of the input waveguides, transforms into many modes in the free space region and then exits through the output waveguides. Another method for making star couplers is to cascade 3 dB couplers, as shown on the right side of Figure 1-26. This works well in planar technologies but it is difficult to assemble fiber based 3 dB couplers because the fibers can not be bent to tight radii.

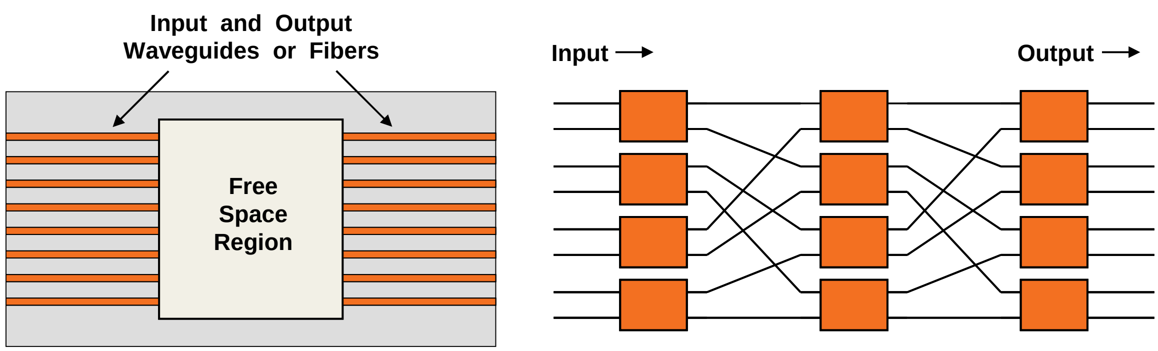

Figure 1-26: Star Couplers Constructed With Free Space Region In Either Fiber or Planar Technologies (Left) and With Cascaded 3 dB Couplers (Right). (40317)

1.6.5 Array Waveguide Grating (AWG)

Diffraction gratings reflect or refract light by different amounts depending on the wavelength of the light. In fiber optics, they are used for wavelength selective functions in external cavity tunable lasers and monochromators and also as muxes and demuxes in WDM systems. One of the most important gratings used for muxes and demuxes in WDM systems is the array waveguide grating (AWG). An AWG can have a single input waveguide with multiple wavelengths on it and separate them onto different output waveguides, or it can take several input waveguides with different wavelengths and combine them onto one output waveguide. An AWG can also be connected as an add-drop multiplexer in which a single channel can be added or removed without affecting other channels on the waveguide. AWGs with up to 128 channels with channel spacing as small as 25 GHz are available.

AWGs are built using planar waveguide technology because of the tight dimensional control required. Figure 1-27 shows an AWG with a single input waveguide and three output waveguides. The input is a star coupler which takes the light from the input waveguide and evenly distributes it to all of the “grating” region waveguides. The lengths of these waveguides differ by a fixed delta which introduces a phase difference at the input to the second star coupler. This light then travels through the free-space region of the second star coupler and arrives at the output waveguides. Due to differences in path lengths through the “grating” region and the second star coupler, light arriving from the different “grating” waveguides will have different phases resulting in constructive and destructive interference at the outputs of the star coupler. The AWG is designed so that for each wavelength, constructive interference occurs at the input to one waveguide and nowhere else. For the orange light in the example in Figure 1-27, constructive interference occurs at the bottom output waveguide.

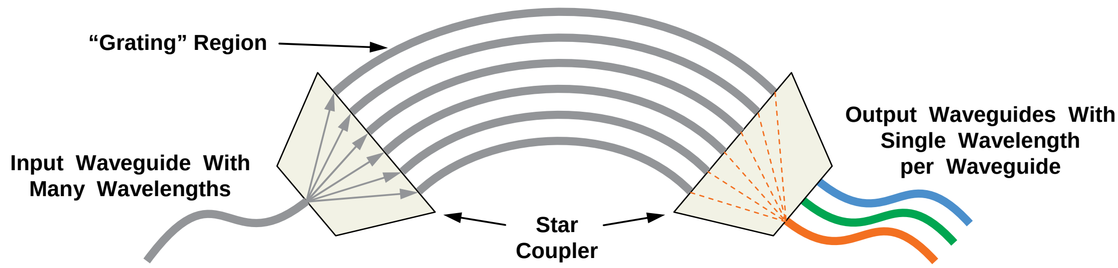

Figure 1-27: Diagram of An Array Waveguide Grating Consisting of Input and Output Fibers, Two Star Couplers and A Grating Region. (40318)

1.7 Optical Switching

To construct a communication network beyond a simple point-to-point network, optical switching elements are required. Early applications of optical switches in the telecommunications industry were to provide route diversity to protect against network failure. These switches are generally opto-mechanical and have relatively slow switching times (on the order of 5 – 10 ms). For modern WDM networks, switches based on the same principles as the modulators and couplers in Sections 3.3 and 1.6.3, respectively, are being used.

Simple switches (see Figure 1-28) include the 1x2 switch based on a Y-coupler and the 2x2 switch. The 1x2 switch is a Y-coupler with control electrodes added to direct the light to one output port or the other. These switches are typically fabricated with electro-optic material, such as lithium niobate, so the electric field from the electrodes changes the RI of the material. The 2x2 switch has two states; the bar state in which the two inputs pass straight through and the cross state in which each input is directed to the opposite output. These switches are constructed using resonant couplers or Mach-Zehnder interferometers on an electro-optic material such as lithium niobate. Mach-Zehnder interferometers are discussed in Section 3.3. We have seen the 2x2 switch referred to as an optical cross-connect switch (OXC) but this term is also used in the telecommunications industry referring to a larger switch that cross-connects separate wavelengths between different fibers. To avoid confusion, we will only use the term 2x2 switch. These simple switches can be used to build more complicated components, such as reconfigurable optical add/drop muxes.

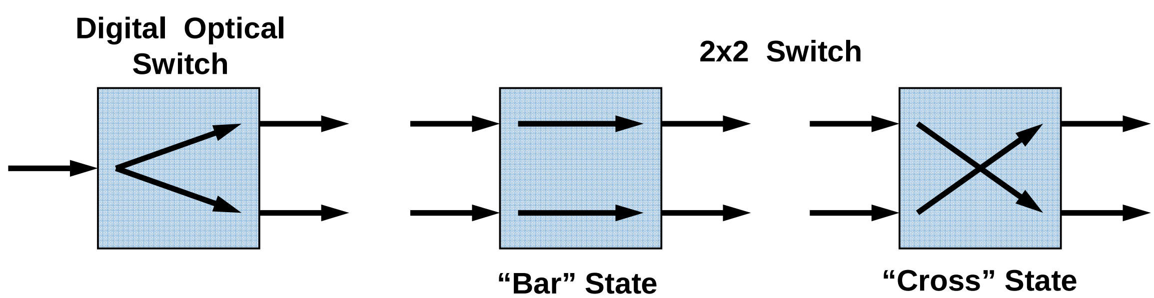

Figure 1-28: Diagram Showing Functions of the Digital Optical Switch and the 2x2 Switch. (40319)

Optical add/drop muxes (OADMs) are used to remove one or more wavelengths from a fiber and to add signals at the same wavelengths to the fiber without affecting the other channels on the fiber. There are two varieties of OADMs: fixed and reconfigurable (ROADMs). Examples of these are shown in Figure 1-29. The attenuators are used to balance the power in all of the channels. In the fixed OADM example, two wavelengths are dropped and added while three are passed through the OADM. For the ROADM, any of the five wavelengths could be dropped and added. The mux and demux can be made using a single AWG and the 2x2 switches can be fabricated on the same substrate to make a ROADM (without attenuators) on a photonic integrated circuit.

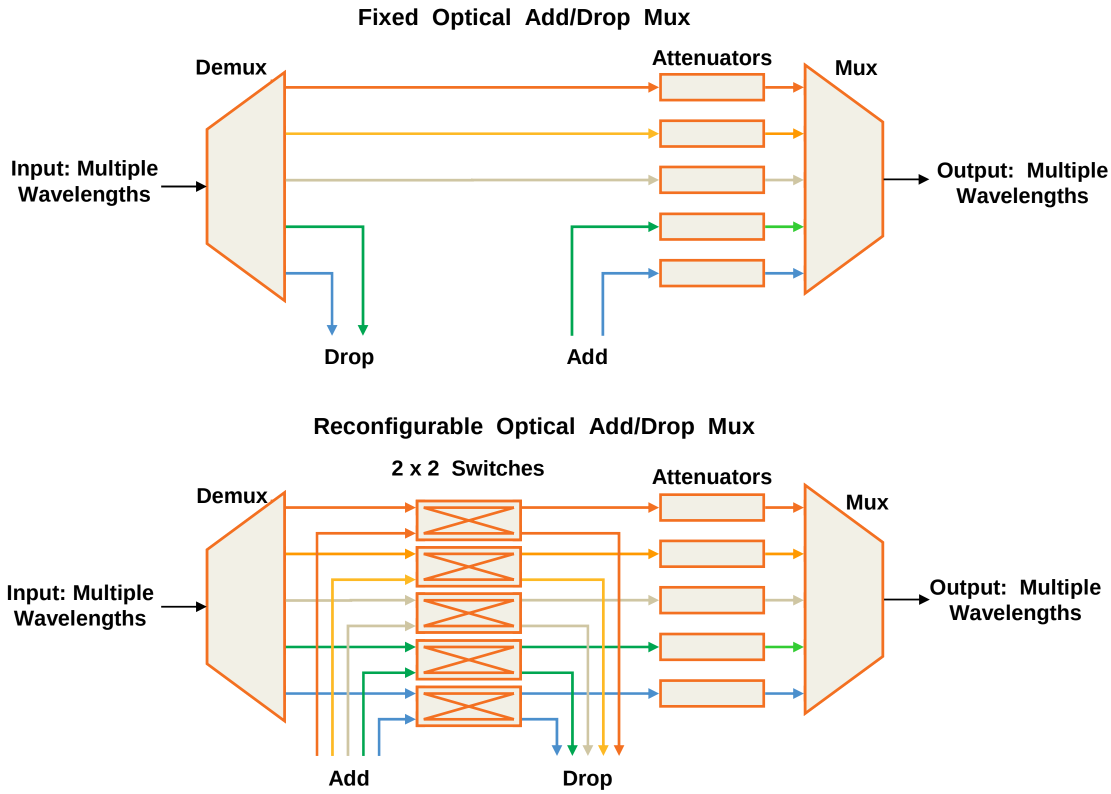

Figure 1-29: Diagram of A Fixed OADM and A Reconfigurable OADM. (40320)

1.8 References

[1-1] Dutton, H.J.R.: Understanding Optical Communications, IBM Redbook 1998. http://www.redbooks.ibm.com/redbooks/pdfs/sg245230.pdf

[1-2] http://en.wikipedia.org/wiki/Fiber_optic_communication (October 2008)

[1-3] Application Note: HFAN-2.2.0 Rev. 1; 04/08, Extinction Ratio and Power Penalty, Maxim Integrated Products. http://pdfserv.maxim-ic.com/en/an/AN596.pdf

[1-4] Sze, S.M.: Physics of Semiconductor Devices, 2nd Edition. John Wiley and Sons, New York 1981.

2 Fiber Types and Other Transmission Media

2.1 Optical Fibers

It is practical to guide light (not necessarily in the visible wavelength range) over long distances using a dielectric fiber of circular cross section, provided that: 1) the fiber material is highly transparent over the intended wavelength range of operation, 2) can be uniformly drawn into fiber form, 3) is mechanically robust, and 4) that its index of refraction can be accurately controlled as a function of radius. For the core material, the index of refraction must be slightly higher than that of the cladding and typically differs by only a few tenths of one percent. Silicate-based glass (basically SiO_2, or fused quartz, with spatially-controlled dopants to modify the index of refraction in the radial direction) and polymers (plastics) are the most common materials used for this purpose. Silicate-based glass fibers are the lowest loss and least dispersive fibers. In general, losses in silica-based glass fibers are very small in the approximate wavelength range of 800 to 1600 nm (near infrared), and are typically on the order of tenths of a dB per km.

Although there are numerous specialty fibers that have been developed for the telecommunications industry, three of the basic types of optical fibers that are commonly used are described here: single-mode (SM), multimode (MM), and polarization maintaining (PM) single mode fiber. SM fiber operating in the 1300 nm to 1550 nm wavelength range is predominantly used in long haul communications since it is least dispersive and most transparent in this range. At 1550 nm wavelength the optical power loss in this type of fiber is only 0.15 dB/km. All long distance underwater cables employ SM fiber and operate at this wavelength [2-1], [2-2]. To achieve single mode propagation, however, requires a small core diameter, in the range of 8 to 10 μm. The cladding layer is typically 125 μm with the index of refraction changing abruptly (step-index) at the interface between the cladding and the core. Figure 2-1 shows actual photomicrographs of the cross-section of these three types of fiber [2-3] with the core illuminated. In all three cases the cladding diameter is 125 μm.

Though SM fiber supports a single mode with linear polarization, there are two possible independent (degenerate and orthogonal) polarizations that can coexist. With a perfectly isotropic and circular fiber core these two polarization modes would have exactly the same propagation characteristics, and a linearly polarized wave launched at the transmitting end would end up linearly polarized at the receiving end. However, for physical fiber there will always be small perturbations in the fiber core material that will result in slight birefringence for the two polarization modes, generally resulting in elliptical and unpredictable polarization at the receiving end. For some applications it is important to ensure linear polarization throughout the fiber length. PM single mode fiber accomplishes this by essentially suppressing one of the possible two polarizations by purposely inducing stress in the core material. Several methods are utilized to induce this stress including embedding special rods within the cladding (PANDA fiber). In Figure 2-1 note the two stress rods in the PM SM fiber embedded in the cladding (light circular shadows – the reason for the name "PANDA").

In contrast to the SM fiber, the core diameter of MM fiber is in the range of 50 to 90 μm (standard core diameters are 50, 62.5 and 85 μm). Historically, before the advent of the erbium doped fiber amplifier (EDFA) in the late 80's, telecommunication systems based on MM fiber were seriously pursued by AT&T, for example [2-1]. One of the advantages of MM is that because of the larger core diameter, alignment to other optical components is less demanding than it is for SM fiber. This is an important consideration from a sub-assembly manufacturing and cost perspective. However, because of the larger core diameter, light propagation is supported simultaneously by multiple modes, not all of which have the same propagation constants. For a 50 μm MM step-index fiber, approximately 200 modes are involved in guiding the light along the fiber. Consequently, in a broadband modulated signal such as an NRZ bit stream with fast rising and falling edges, the well defined pulses at the sending end will progressively be distorted (pulse spread) as they progress down the fiber because of modal dispersion. Compensation techniques to mitigate modal dispersion involving special index of refraction profiles have evolved. Specifically, graded-index fiber employs a nearly quadratic dependence of the index of refraction as a function of radius within the core, having the highest value at the center of the core. The number of modes are reduced by a factor of two compared to the step-index profile, and in theory modal dispersion can be reduced by three orders of magnitude [2-3]. Fabrication conditions are naturally more demanding for this type of fiber and the theoretical limit for modal dispersion reduction can only be approached. Nevertheless, MM graded index fiber can support communications up to around 100 Mbps for distances up to 100 km [2-4].

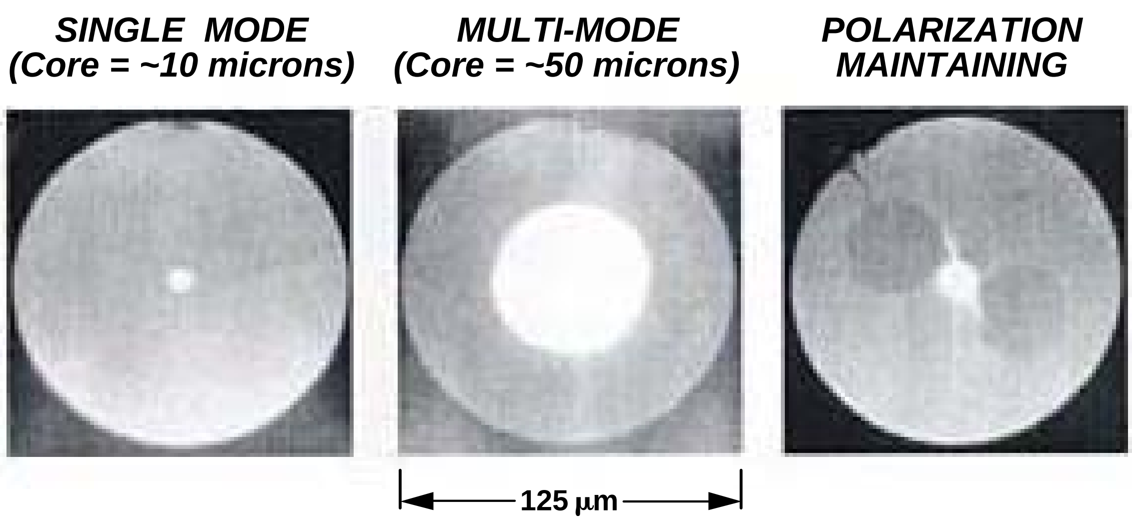

Figure 2-1: Three Basic Types of Optical Fibers: (a) SM – Core About 10 μm, (b) MM – Core Around 50 μm, and (c) SM PM Fiber; Cladding Diameter is 125 μm for All Three Cases. (40257)

2.2 Polymer Fibers and Optical Waveguides

Instead of glass, plastic fibers made of high quality polymers can also be used to guide light over relatively short distances. The light can be from visible through to the near-infrared wavelengths. Attenuation can be below 0.001 dB/cm (100 dB/km), which is much too high for telecom applications but is acceptable for distances up to approximately 300 m, and bandwidth-distance products around 100 MHz-km are feasible for multimode fiber [2-3],[2-5]. A popular material for this application is polymethyl methacrylate (PMMA), commonly known as plexiglass. Other plastics suitable for optical fiber include polystyrene and polycarbonate. Polymer fibers offer lower cost and greater mechanical flexibility (smaller bend radius) than their silicate glass counterparts. Although step-index multimode fibers are the most common polymer fibers in use, single mode fiber and graded index multimode fiber is also available.

Polymer optical waveguides, (not necessarily of circular cross section) can also be batch-fabricated and embedded in other plastics for use in short-reach communications links, such as might be used for high performance compute (HPC) systems. A recent research project, described in reference [2-6], demonstrated the use of multiple polymer waveguides to form dense parallel optical communication channels. An objective of the project was to demonstrate a potentially cost effective, high bandwidth optical link with dense packing of parallel polymer waveguides embedded in standard FR4 printed circuit board material. Called an “optoboard,” this array of optical waveguides was used in a complete operational high speed data link incorporating 16 bidirectional channels, with each channel supporting up to 15 Gbps data rate. The optoboard was approximately 15 cm long between transceivers and the embedded polymer waveguides were laminated onto the FR4 board on a 62.5 μm pitch. The optical core dimensions were 35μm x 35 μm. An example of this technology is shown in a more recent result from IBM’s Zurich Research Laboratories [2-7] in Figure 2-2. In this case a two dimensional 4 x 12 array of polymer waveguides with a 250 μm pitch in both the x- and y-directions is shown. With core dimensions of approximately 50 μm, the optical loss at 850 nm is only 0.028 dB/cm [2-7]. With a demonstrated bit rate of 12.5 Gbps per channel, the aggregate bit rate density is potentially 2 Tbps in a 1 cm length of circuit board!

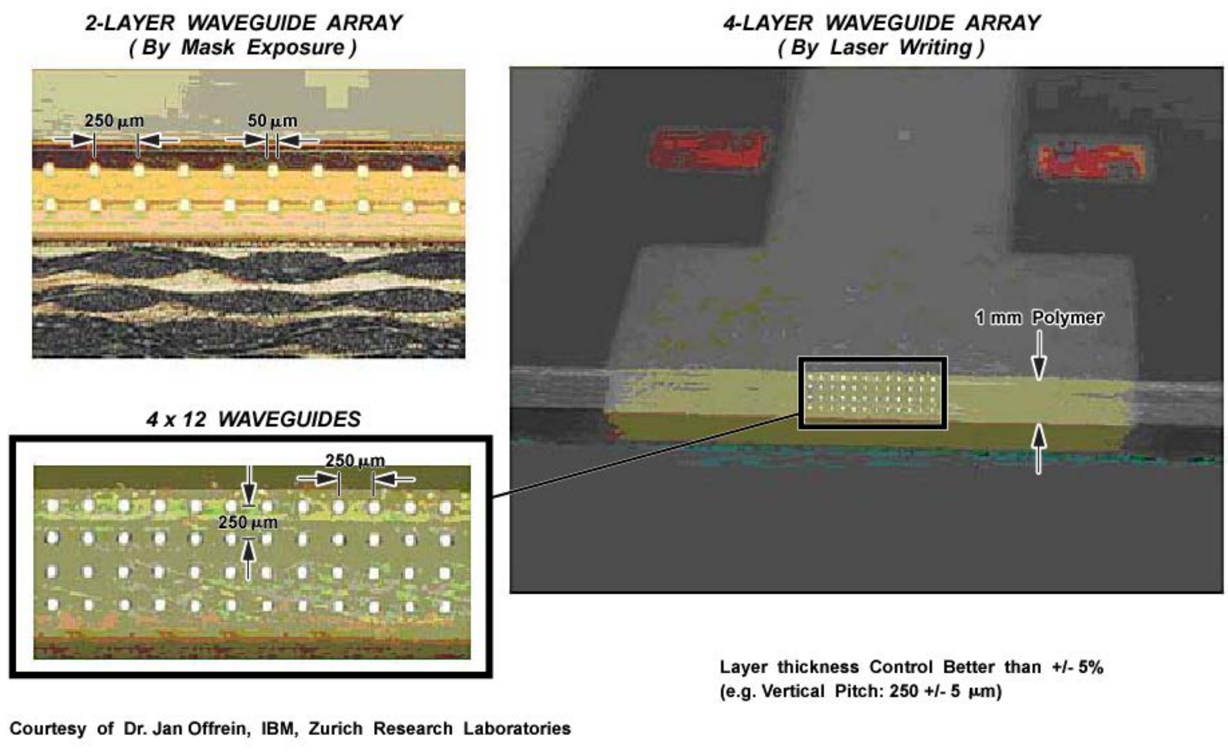

Figure 2-2: A Research Example of Multiple Polymer Optical Waveguides. (40258)

2.3 Optical Waveguides for Planar Lightwave Circuits (PLC)

For optical hybrid integrated circuits (i.e., an assembly comprised of a suitable substrate onto which individual optoelectronic chips are mounted), or monolithic photonic integrated circuits (PICs), dielectric optical interconnects are required between chips or between devices on the same chip. Using separate optical fibers for interconnects is generally impractical because of the optical and mechanical incompatibility of the interface between fiber and optical device. Instead, it is possible to fabricate miniature optical waveguides either on (or within) the substrate, or directly on the monolithic circuit itself using photolithographic techniques similar to those used in semiconductor IC fabrication. The substrate, or semiconductor, must therefore have optical properties consistent with the intended waveguide formation. Examples of suitable substrates include a silica glass wafer, a silicon or III-V semiconductor wafer, a block of plastic, a crystal of lithium niobate, among others. In all cases the fabrication process defines the optical waveguide through variation of the index of refraction in the cross section of the guide i.e., transverse to the propagation direction.

Figure 2-3 shows the concept of an optical hybrid IC. The substrate serves as a conventional circuit board for the electrical interconnects, and as a mechanically stable platform for mounting the various optoelectronic integrated circuits (OEICs). It also is the base material for fabricating the optical dielectric waveguides. Besides their use as optical interconnects between OEICs, the dielectric waveguides themselves can be part of a variety of planar subcomponents, such as waveguide bends, directional couplers, waveguide crossovers, power splitters, etc. In order to guide light, however, channel waveguides of small dimensions must be formed in the substrate that are highly transparent and uniform in the direction of propagation, yet confine the light to the channel by means analogous to the optical fiber, i.e., through appropriate variation in the index of refraction in the transverse direction.

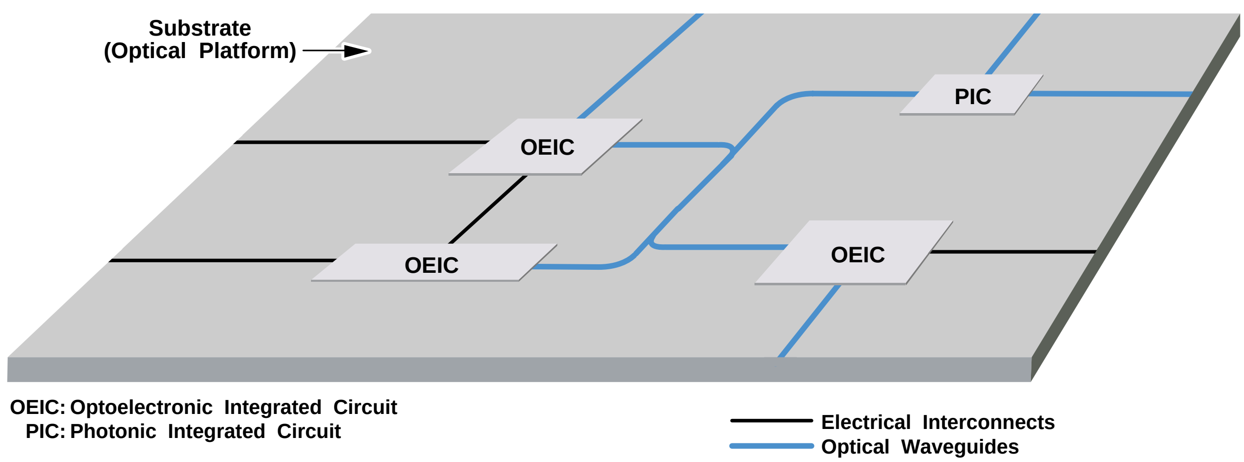

Figure 2-3: Diagram Illustrating Basic Optical Hybrid Integrated Circuit. (40236)

Several basic examples of channel waveguides are shown in Figure 2-4 [2-8],[2-9]. The rib and ridge waveguides are formed at the substrate's surface, while the other two examples show the waveguide core embedded within the substrate. In all cases the core material, in which the light is concentrated, has the largest index of refraction relative to the other surrounding regions and layers. Generally the waveguides are designed to operate in single mode since many of the optical subcomponents require it (for example, a Mach-Zehnder device relies on constructive (destructive) interference of the optical fields). For monolithic implementations, typical waveguide widths are in the range of 2 to 5 μm with losses less than about 1 dB/cm [2-3].

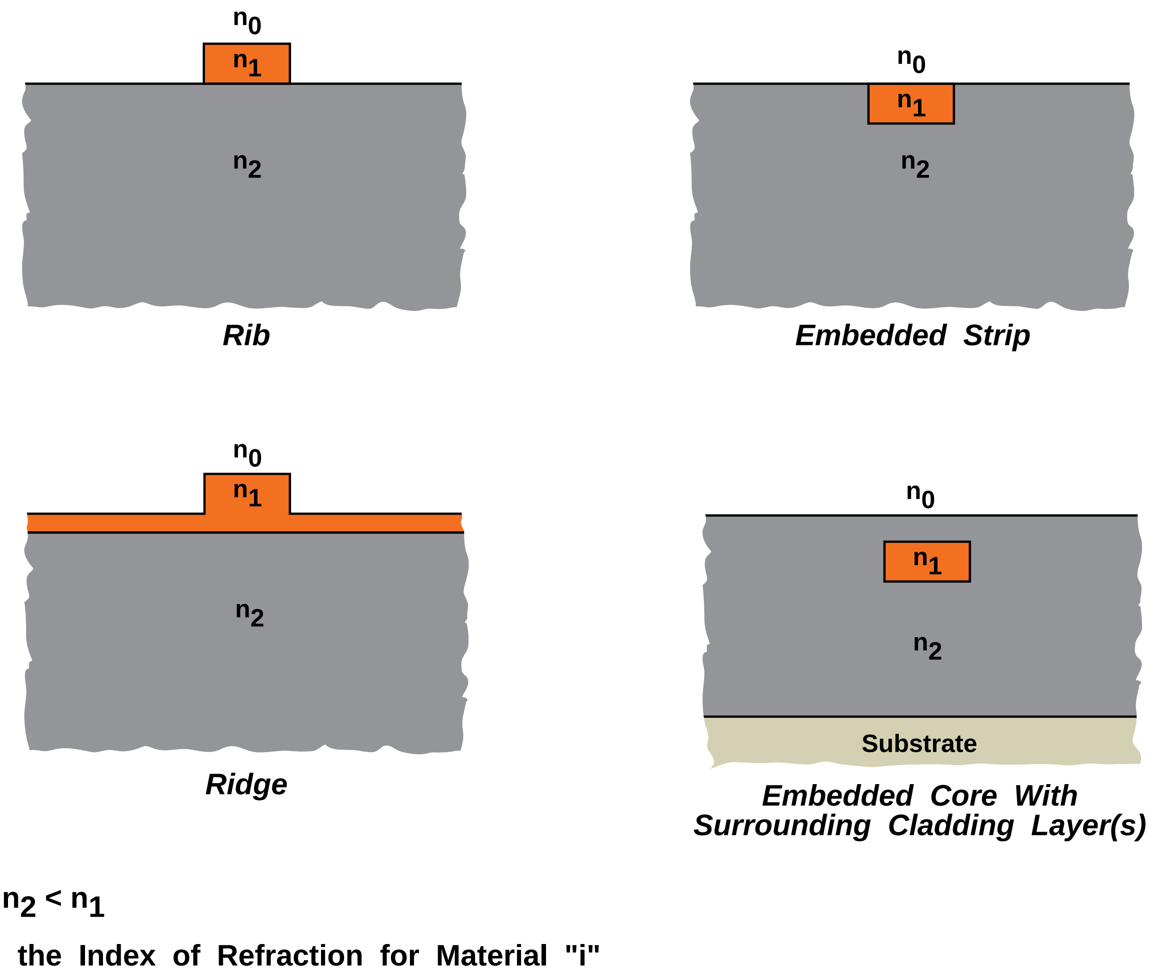

Figure 2-4: Several Types of Optical Waveguides That Can Be Fabricated on Planar Substrates. (40240)

An important practical system for a hybrid optical IC structure is the SiO_2/Si system depicted in Figure 2-5 [2-3],[2-10]. The optical platform in this case is a silicon substrate with etched pedestals for flip-chip mounting the OEICs. The optical waveguide, core and cladding, is formed in silica (SiO_2) using the flame hydrolysis process. Etch stop techniques used in the pedestal fabrication process ensures precise vertical alignment between the optical waveguide and the optical port of the OEIC. The pedestal also provides an excellent thermal path from the active device to the silicon substrate. Excellent low loss performance of 0.1 dB/cm is achieved for the optical silica waveguides. For electrical connection, the electrical lines reside on top of the SiO_2 layer and thus have much less capacitive loading than they would have if they were fabricated on top of the silicon substrate.

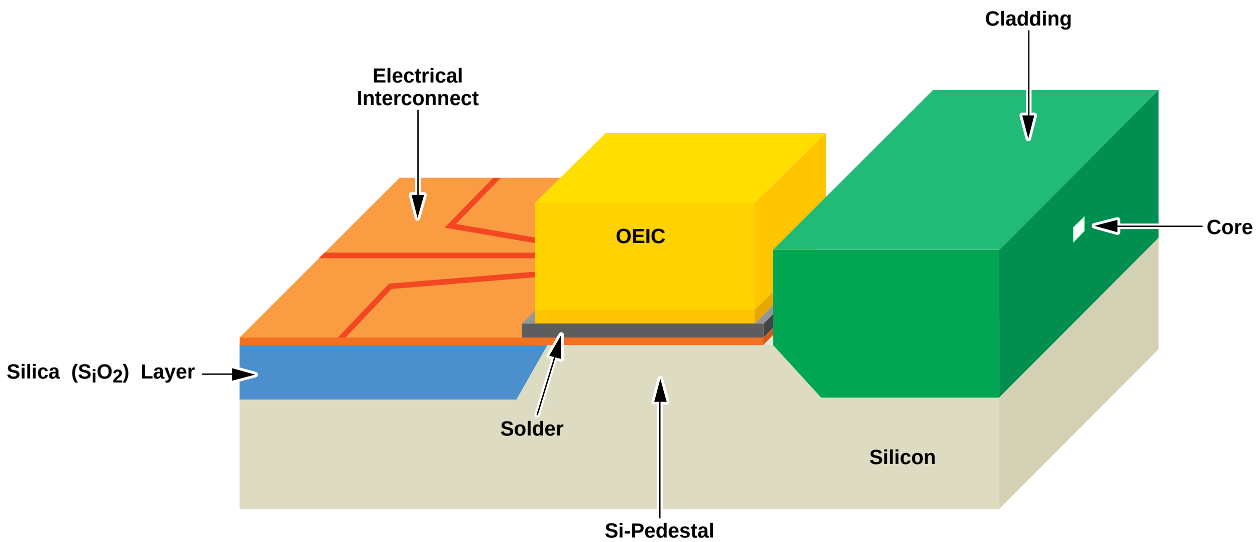

Based on Concept First Described in Reference 2-8

Figure 2-5: Diagram of Optical Hybrid IC Implemented in Silicon. (40245)

In general, fabrication of such hybrid optical integrated circuits is a topic in itself and is beyond the scope of this reference guide. Suffice it to say that precision semiconductor fabrication techniques such as lithography, epitaxial growth, etching, and doping, among others, are employed to form these types of optical subassemblies in Si, as well as in other optically suitable substrates.

2.4 References

[2-1] Gambling, W. A.: The Rise and Rise of Optical Fibers. IEEE Journal of Selected Topics, 6(6):1084-1093 (November/December) 2000.

[2-2] Bergano, N. S.: Wavelength Division Multiplexing in Long-Haul Transoceanic Transmission Systems. Journal of Lightwave Technology, 23(12):4125-4139 (December) 2005.

[2-3] Cox, C. H.: Analog Optical Links. Cambridge University Press, 2004. ISBN 0-521-62163-1.

[2-4] Grote N., and H. Venghaus (Editors): Fibre Optic Communication Devices. Springer-Verlag, New York, 2001, p. 37. ISBN-13: 9783540669777.

[2-5] Paradigm Optics, Inc., Web Ad: http://www.paradigmoptics.com/pof/poffaqs.html

[2-6] Doany, F. E., et. al.: Terabus: A 160-Gb/s Bidirectional Board-Level Optical Data Bus. IEEE LEOS The Society for Photonics, Conference Proceedings of the 20th Annual Meeting of the IEEE Lasers and Electro-Optics Society, pp. 545-546, October 21-25, 2007, Lake Buena Vista, FL.

[2-7] Private communication, courtesy of Dr. Bert Offrein, Zurich Research Laboratory.

[2-8] Das, P.: Lasers and Optical Engineering, Springer-Verlag, New York, 1990.

[2-9] Dutton, H. J. R.: Understanding Optical Communications. IBM International Technical Support Organization, p.150, September 1998. http://www.redbooks.ibm.com

[2-10] Yamada, Y., et al.: Silica-Based Optical Waveguide on Terraced Silicon Substrate as Hybrid Integration Platform. Electronics Letters, 29(5):444-446 (March) 1993.

3 Light Source Options

There are a variety of lasers and light emitting diodes (LEDs) available as light sources for fiber optic communications. Many factors go into choosing which is best for a given application. These factors include economics, laser output power, link length, bit rate, link loss and power dissipation. In this section, we will discuss trade-offs between different types of lasers and LEDs and describe which of these light sources are the most promising for high performance compute (HPC) applications. We will also discuss modulators and optical amplifiers in this section.

3.1 Light Emitting Diodes (LEDs)

LEDs have been used in communication systems at 850 nm and 1310 nm. Historically they have been less expensive than laser diodes but that has changed as LED structures have become more complicated and as laser diode production has increased. The main problem with using LEDs for HPC applications is modulation speed. They can be directly modulated only to several 100 MHz which is far too slow and adding an external modulator increases cost and complexity. We have not seen any references to LEDs with external modulators.

3.2 Lasers

The basic operating principles of lasers were introduced in Section 1. Here, we will discuss the different types of semiconductor laser diodes in greater detail.

3.2.1 Fabry-Perot (FP) Lasers

Fabry-Perot (FP) lasers are the standard edge-emitting lasers with mirrors on the cleaved edges of the semiconductor chip discussed in Section 1. The basic FP laser is just an LED with mirrors on two opposite ends of the device. The output wavelength is determined by the bandgap of the material used and by the length of the cavity. As discussed in Section 1, semiconductor LEDs and laser diodes emit light over a range of wavelengths due to the energy distribution of electrons and holes in the semiconductor. The length determines which modes or wavelengths are resonant in the Fabry-Perot cavity. Modes in a laser are not the same concept as modes in a fiber or optical waveguide and should not be confused with them. The laser modes are determined by the equation

$$\lambda = 2nL/x,$$

where λ = wavelength, n=refractive index of the semiconductor, L= cavity length and x=integer mode index. The mode index is the number of wavelengths that fit in twice the cavity length. For example, with a length L=200 μm and refractive index n=3.45, the mode index x extends from one to infinity and the wavelengths λ go from 2nL (1.38e6 nm) down to zero. FP lasers are typically several hundred microns long so for wavelengths near 1550 nm, the mode index is approximately 900 and the resonant wavelengths are spaced by approximately 1-3 nm. The spectral characteristics of the laser diode are then determined by the overlap of the emission characteristics (gain spectrum) of the semiconductor and the modes of the cavity as shown in Figure 3-1. The spectral width of the laser is typically given as the full width at half maximum (FWHM) amplitude as shown in Figure 3-1. The individual modes in the laser output are referred to as lines and the linewidth is also measured at FWHM. The basic FP laser produces approximately seven lines, while more complicated designs produce from one to three lines [3-1].

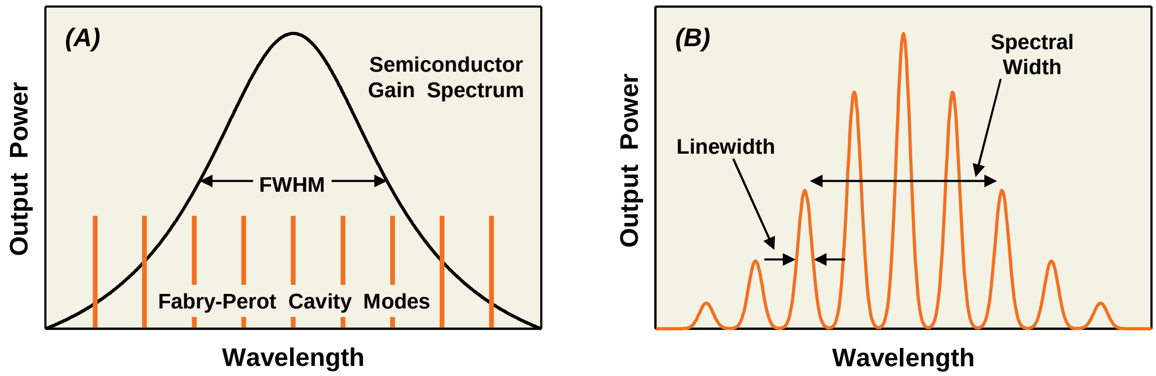

Figure 3-1: *(A)* Idealized Semiconductor Gain Spectrum and Fabry-Perot Cavity Modes of A Laser Diode; *(B)* Idealized Spectral Characteristics of A Laser Diode. (40533)

Typically the light from any semiconductor laser can be launched into single-mode fibers or multimode fibers or waveguides. When lasers with more than one line are coupled to single-mode fiber, each individual line is launched into the fundamental mode for its wavelength. This is similar to a wavelength division multiplexing (WDM) system except the laser lines are very close together in wavelength.

An important laser characteristic is the threshold current, which is the amount of current needed for the laser to start lasing. Figure 3-2 shows a typical optical power versus drive current diagram. Below threshold, the emission is predominantly spontaneous, while above threshold it is mostly stimulated emission. Figure 3-3*(A)* shows a drawing of an unguided FP laser. The vertical black lines represent the current flow through the device and the red lines represent the laser light being emitted from the end of the laser. As can be seen, the current and light emitted are uniformly distributed across the width of the laser. One technique used to reduce the threshold current (an important consideration for reduced power consumption) and also the number of modes in the laser is to use an electrical contact that is smaller than the width of the laser as shown in Figure 3-3*(B)*. The black lines again represent the current flow, which in this case is localized under the contact. Only the active region directly under the contact lases because this is the only region that has enough gain (current) to reach the threshold current. The light output power is reduced because a smaller region of the device is lasing. The electrons and holes in the active region increase the refractive index (RI) of this region by a small amount resulting in some guiding of the laser light along the length of the laser. To increase this guiding, the higher RI active region is embedded in a lower RI cladding region as shown in Figure 3-3*(C)*. This results in somewhat lower light output power but reduces the number of modes to between one and five and reduces the spectral width. This device is referred to as an index guided FP laser. The spectral characteristics of gain and index guided FP lasers are compared in the table in Figure 3-3*(D)*.

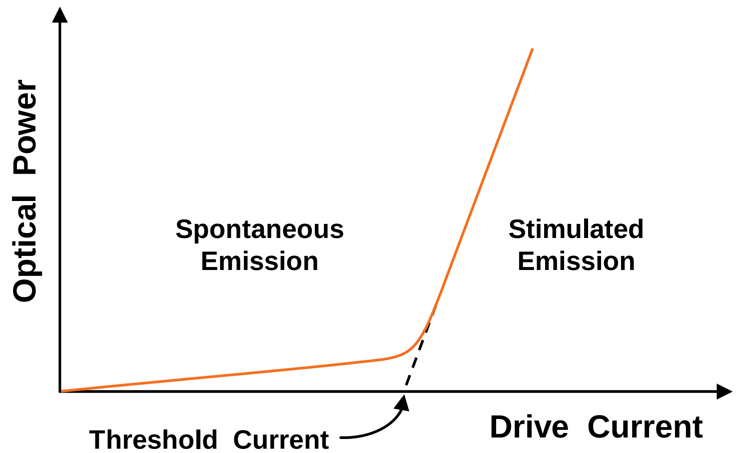

Figure 3-2: Typical Optical Power Versus Drive Current Diagram Showing the Threshold Current of A Laser Diode. (40534)

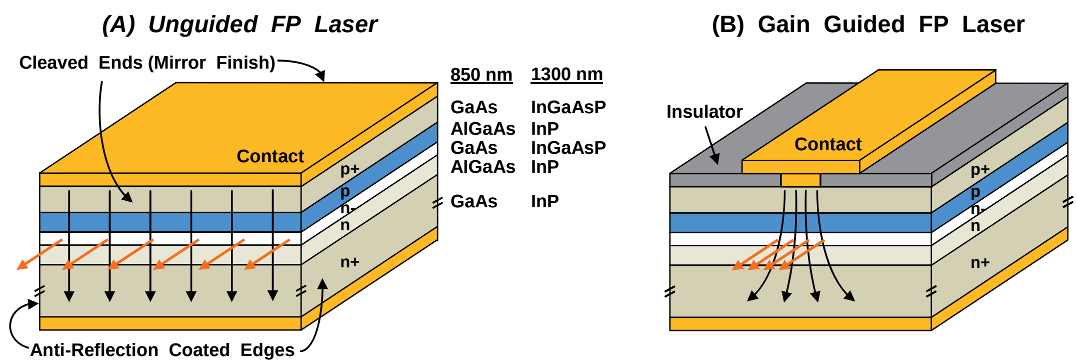

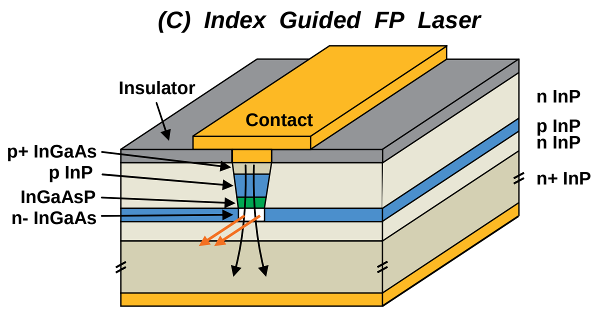

(D) Comparison of Gain and Index Guided FP Lasers

	Spectral Width (nm)	# Lines	Line-Width (nm)
Gain Guided	5-8	8-20	0.005
Index Guided	1-3	1-5	0.001

Figure 3-3: *(A)-(C)* Different Guiding Techniques Used In Fabry-Perot Lasers and *(D)* Comparison of Spectral Characteristics of Typical Gain and Index Guided FP Lasers. (40535)

A few other laser properties, which are generally sources of noise, are described next. These descriptions are taken, nearly verbatim, from [3-1].

After a very short time in operation lasing tends to use up the available excited electrons in the center (dominant mode path) of the cavity. This happens because it is difficult to get power to the entire active region at an even rate. Thus a “hole” is burned in the path taken by the dominant mode. The dominant mode is thereby significantly reduced in power.

Hole burning causes mode hopping. When the strong, dominant mode decreases other modes are able to increase and become dominant. Thus the laser produces light in one mode for a very short time and then it “hops” to another mode, and then to another and then to another. The whole range of resonant modes within the gain spectrum may be covered. This happens very quickly (a few tens of picoseconds per hop) so the laser essentially produces a band of wavelengths. It is important to note that when a single mode dominates the others are usually not suppressed entirely but they are strongly attenuated.

When the signal is sent on a dispersive medium mode hopping can become an additional source of noise. This is because each mode is at a different wavelength and each wavelength will travel at a different speed within the fiber. Not only will the pulse disperse but the dispersion will be irregular and random in nature.

Immediately after power is applied to a laser there is an abrupt change in the carrier (electron and hole) flux density in the cavity caused by the lasing operation itself. This density of charge carriers is one factor that affects the refractive index. In addition, the temperature in the cavity increases quite rapidly. This temperature increase is too localized to affect the length of the cavity immediately but it does contribute to changing the refractive index of the material in the active region.

These changes in the RI of the cavity produce a rapid change in the center wavelength of the signal produced. In the case of semiconductor lasers a “downward” chirp is produced. The wavelength shifts to a longer wavelength than it was immediately at the start of the pulse. It is not a large problem in short distance single-channel transmissions but in long distance applications and in WDM systems, chirp can be a very serious problem. This is due to the fact that it broadens the spectral width of the signal and interacts with other aspects of the transmission system to produce distortion. Indeed, the chirp problem is the main reason that people use external modulators for transmission rates in excess of 1 Gbps.

When the laser is turned on there are short term fluctuations in the intensity of the light produced called “relaxation oscillations”. When power is applied to the laser the upper energy state population builds up until an inversion occurs and lasing can commence. However, lasing can deplete the upper energy state very quickly and if pumping isn't quite fast enough lasing will momentarily stop. Very soon afterwards it will start again as the pump builds up a population inversion again.

This effect varies widely between types of laser. Some can turn on with little or no relaxation oscillation; others (if the pump is a little weak for example) can produce these oscillations continually and never reach a stable lasing state. Most semiconductor communications lasers produce some relaxation oscillation at the beginning of each pulse but stabilize quite quickly.

Relative intensity noise (RIN) refers to a random intensity fluctuation in the output of a laser. The primary cause here seems to be the random nature of spontaneous emissions. As the laser operates new spontaneous emissions occur and some of them can resonate within the cavity and are amplified. This causes some fluctuation in output power.

The random changes in emissions which cause RIN are by nature different in phase from previous emissions. This causes random changes in phase during laser operation resulting in phase noise. This variation is a natural consequence of the way lasers operate and cannot be suppressed. However, the effect is not important in amplitude modulated systems.

Intercavity noise is caused by reflections returning a portion of the optical signal back into the laser cavity. When a signal returns into the laser cavity due to a spurious reflection it is at exactly the right wavelength and will be amplified in the cavity. This causes unwanted fluctuations in the light output.

In general there are two kinds of intercavity noise. The first is caused by nearby reflections such as from the laser-to-fiber coupling. This can be minimized by using anti-reflection coatings and lens couplings designed to minimize reflections. The second is caused by reflections from more distant optical components. In some systems (especially long distance systems) an optical isolator is used immediately following the laser to eliminate the problems caused by these far-end reflections.

After the device has been operating for a while the temperature of the device can change (it will heat up) and this will affect the cavity length and the wavelength will drift (change). Also, there are effects caused by age of the device and the materials it is made from. Some of these effects can cause a slow change in the wavelength over time.

3.2.2 Distributed Feedback (DFB) Laser

The distributed feedback (DFB) laser was developed to further reduce the spectral width of FP lasers. A DFB laser is just an index-guided FP laser with a Bragg grating (a periodic variation of the RI index) embedded in the semiconductor layer next to the active layer, as shown in Figure 3-4. The fields from the light wave extend into this adjacent layer and interact with the grating resulting in reflections at each RI change. Constructive interference occurs when the period of the grating is a multiple of the wavelength of the light and other wavelengths interfere destructively and are not reflected. To avoid having FP modes present in the laser cavity, one end of the laser has a mirror finish to reflect light back into the laser and the other has an anti-reflection coating to allow light to exit the laser. DFB lasers typically have very narrow linewidths (~0.0002 nm), low chirp and very low RIN. However, they are also 1) very sensitive to reflections which disturb the laser's stable resonance, 2) sensitive to temperature, and 3) relatively expensive.

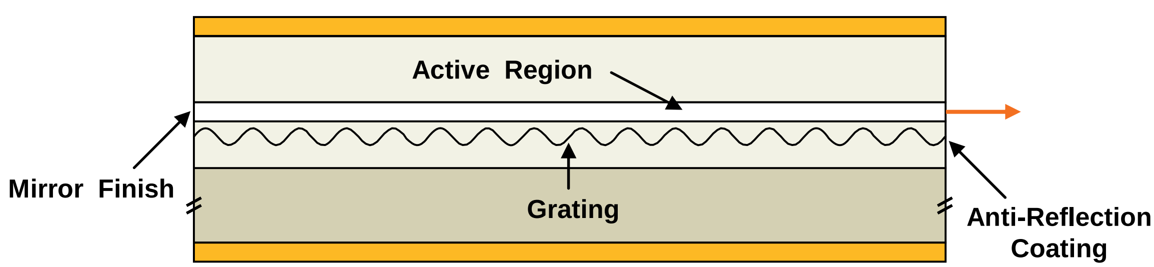

Figure 3-4: Drawing of the Side View of A DFB Laser.   (40536)

3.2.3 Distributed Bragg Reflector (DBR) Laser

The distributed Bragg reflector (DBR) laser is similar to the DFB laser except the Bragg grating is not in the active area of the laser as shown in Figure 3-5. There are two gratings, one at each end of the laser, which act like mirrors reflecting some of the laser light back into the laser cavity. The active region of the laser heats up from current flow through it, resulting in changes of the RI in that region. In the DFB laser, the grating is in the active region and the RI changes with temperature result in changes in the output wavelength of the laser. In the DBR laser, the gratings are not in the active region and so the wavelength is not as temperature sensitive. DBR lasers typically produce a single line only with a linewidth of around 0.0001 nm [3-1]. They are, however, less efficient because of absorption in the gratings.

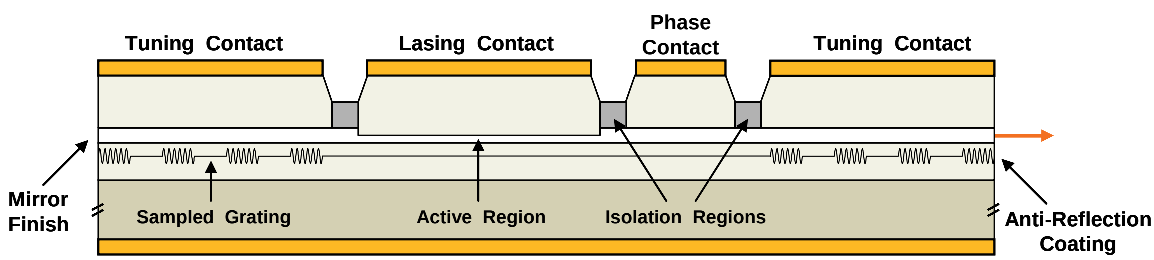

Figure 3-5: Drawing of the Side View of A Tunable DBR Laser With Sampled Gratings. (40537)

In WDM applications, it is necessary for the laser to be tunable over a range of wavelengths. Coarse tuning is used to select a wavelength channel on the International Telecommunications Union (ITU) grid and fine tuning is used to center the wavelength in the channel. Coarse tuning in DBR lasers is accomplished by passing a current through the grating region which causes a change in RI resulting in a different resonant wavelength. This gives discontinuous tuning over approximately 10 nm. To increase the tuning range, sampled gratings are used, as shown in Figure 3-5. The RI in the gratings is tuned so that resonant wavelengths in each grating overlap as illustrated in Figure 3-6. DBR lasers with sampled gratings can be tuned over approximately 100 nm. Fine tuning is done by passing current through a region between the active region and one of the gratings. This region is under the phase contact in Figure 3-5.

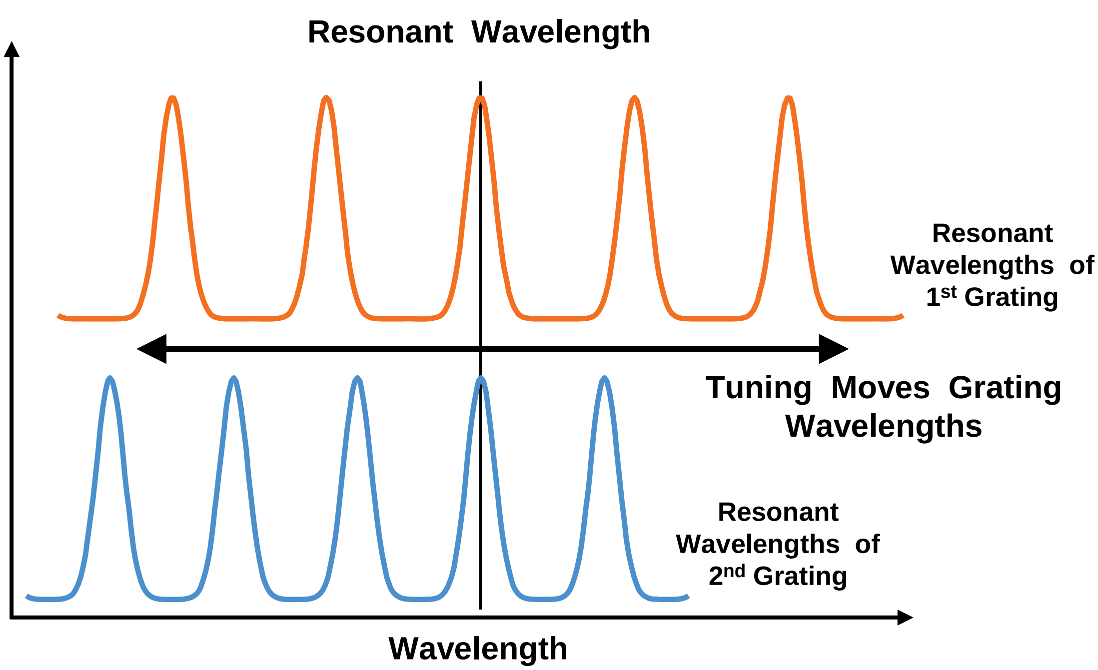

Figure 3-6: Schematic Drawing of the Resonant Wavelengths of Two Sampled Gratings. (40538)

3.2.4 Vertical-Cavity Surface-Emitting Laser (VCSEL)

While the semiconductor edge-emitting laser was first demonstrated in the early 1960s, vertical-cavity surface-emitting laser (VCSEL) technology development is relatively recent. Heightened research interest and laboratory demonstrations of VCSELs were prevalent during the late 1980s. Large scale production and fielded devices, especially GaAs 850 nm VCSELs, occurred towards the end of the 1990s.

The VCSEL takes a DBR laser and turns it on its side. So VCSEL structures have distributed Bragg reflector mirrors above and below the active region as shown in Figure 3-7. This results in a cavity length that is at least two orders of magnitude smaller than in edge emitting lasers. The distributed Bragg reflector mirrors consist of alternating layers of semiconductor with different RI with the materials and layer thicknesses chosen to reflect the desired wavelength of light. The laser light is emitted either through the top surface or the bottom surface depending on the relative reflectivities of the two mirrors. The active region consists of a p-n junction and typically multiple quantum wells. Multiple quantum wells are used to increase the thickness of the active region resulting in increased interaction between the laser light and excited electrons and thus increasing the gain of the laser. The current density in the active region is increased by forming an electrically conductive aperture surrounded by an insulating layer just above the active region as shown in Figure 3-7. The insulating region is formed by proton bombardment or by selective oxidation of a semiconductor material with high aluminum content. The diameter of the aperture can be as small as several microns for low threshold currents (~100 µA) or as large as 100 µm for high output powers (~100 mW) [3-2]. VCSELs can be single-mode for smaller diameter devices (<15 µm; 12 µm typical) or multimode for larger diameter devices (>15 µm; 20 µm typical).

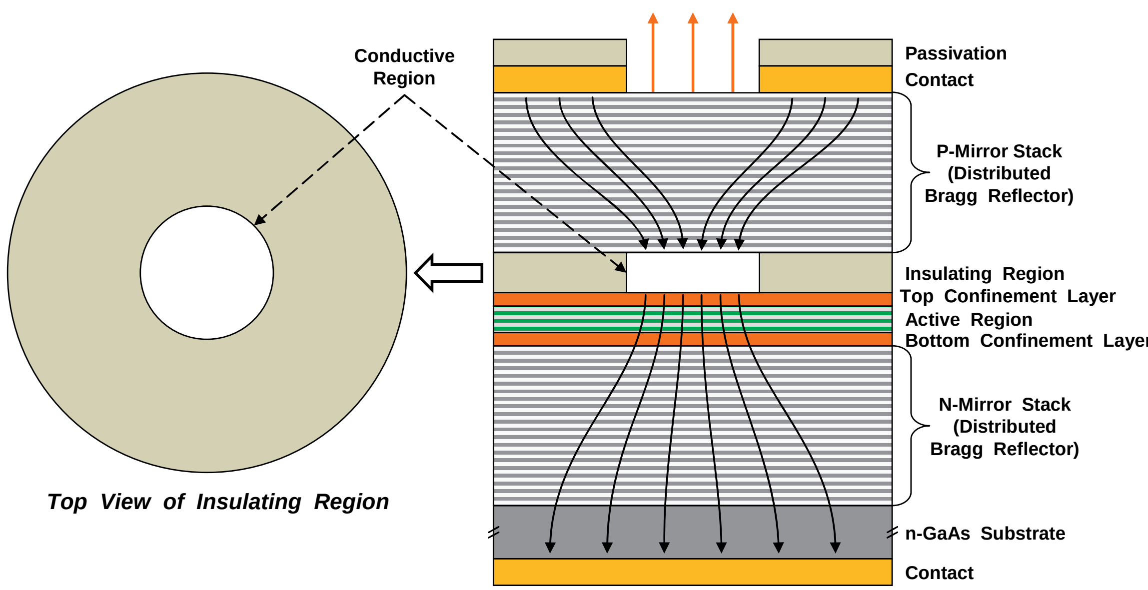

Figure 3-7: Drawing of the Side View of a VCSEL. (40539)

VCSELs have several advantages over edge-emitting lasers [3-2]. These include:

A. Ultralow threshold operation
B. High-speed modulation capability

C. 1D and 2D laser arrays can be formed
D. High-power conversion efficiency
E. Relatively insensitive to temperature variation
F. Easy coupling to optical fibers
G. Initial probe test can be performed before separating devices into discrete chips.

3.3 Modulators

Modulators are made from materials that change their optical properties under different applied fields. The main effects used are:

A. Electro-optic and magneto-optic effects: The material's optical properties (RI or polarization) change in the presence of either an electric or magnetic field.

B. Electro-absorption effect: The material's optical absorption properties change in the presence of an electric field.

C. Acousto-optic effect: Very high frequency sound waves are used to deflect light from one path to another in a crystal.

One of the most common types of optical modulators is the Pockels cell, which is based on an electro-optic effect. Figure 3-8 shows a diagram of a Pockels cell. The input light is polarized in one direction by the first polarizer, then in the electro-optic material the polarization can be rotated by the electric field provided by the voltage applied to the electrical contacts. Finally the light passes through a second polarizer which is rotated 90° from the first polarizer. For a '1,' the polarization is rotated 90° in the electro-optic material so that it passes through the second polarizer, while for a '0,' the polarization is not rotated and the light is absorbed by the second polarizer.

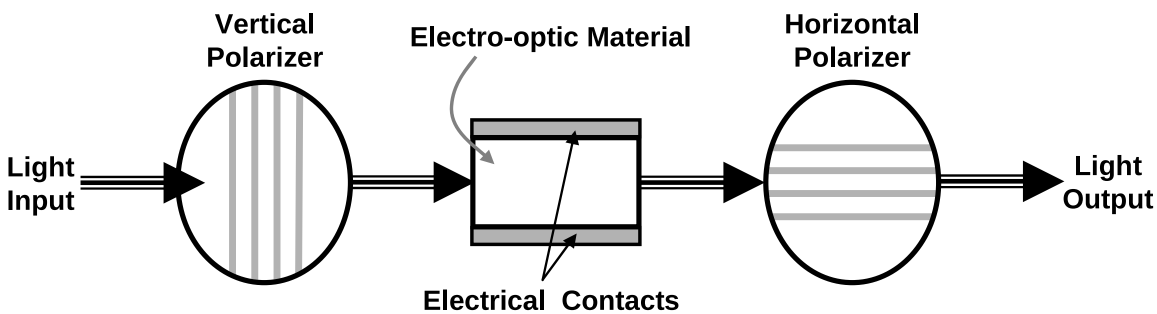

Figure 3-8: Diagram of Pockels Cell Modulator. (40297)

Electro-optic materials that are typically used are lithium niobate ($LiNiO_3$), ADP (ammonia dihydrogen phosphate – $NH_4H_2PO_4$), and KDP (potassium dihydrogen phosphate – KH_2PO_4). The input light must be polarized so these modulators work well at the exit point of lasers. However; they do not work well with standard fiber in which the polarization changes randomly. Another drawback is that high voltages (approximately 1000 V) are needed which can limit the modulation speed and increase the cost. These modulators can operate up to 1 GHz and have extinction ratios from 100 to 1000.

Faraday effect modulators are similar to Pockels cells, except that they employ a magneto-optic effect to rotate the polarization rather than an electro-optic effect. They use an electromagnet to generate and modulate the required magnetic field. These modulators are generally not used in communications because they are slower than Pockels cells and somewhat more expensive.

Another important type of modulator is the electro-absorption modulator (EAM). This device utilizes two separate effects to reduce the bandgap energy of a semiconductor in an electric field. As discussed in Section 1.5.2, light with wavelength less than $1.24/E_g$ will be absorbed while longer wavelengths will pass through the semiconductor. When the bandgap is reduced in the modulator, light with a slightly larger wavelength ($\lambda = 1.24/E_{g'}$) will be absorbed. Thus light with wavelength $\lambda = 1.24/E_{g'}$ can either be absorbed or pass through the modulator depending on the state of the modulator. EAMs use a reverse biased p-n junction to apply the electric field across the semiconductor. When light is absorbed by an EAM, electron-hole pairs are generated, just like in a photodiode. However, the operation of the EAM depends on the electric field and not on the current flow so the modulation speed depends on how fast the electric field can be changed and not on the flow of the charged carriers. Therefore indirect modulation using an EAM can achieve higher data rates than a directly-modulated laser (> 10 Gbps). Discrete EAMs are generally not used because they have relatively high loss (~9-12 dB). However, they can be integrated onto the same chip with a laser and in this configuration the loss is much lower (~1 dB). Figure 3-9 shows an EAM integrated with a semiconductor optical amplifier (SOA) and a sampled grating-distributed Bragg reflector (SG-DBR) laser.

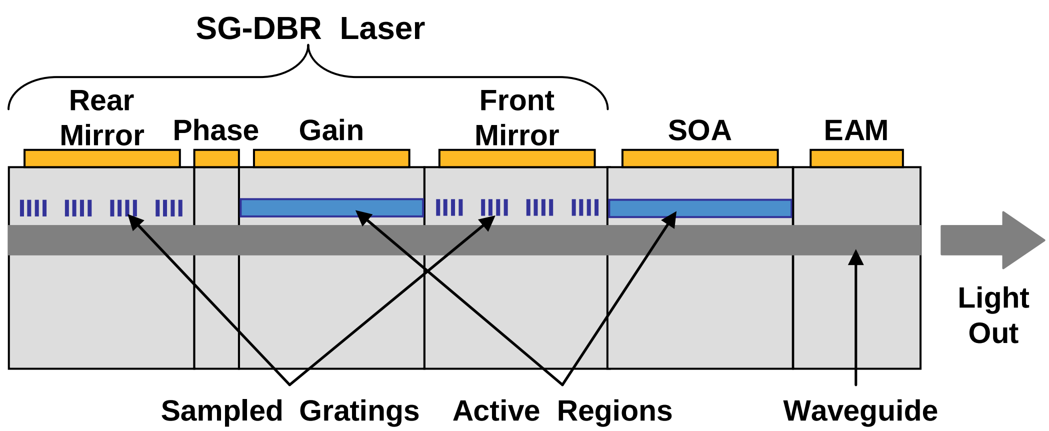

Figure 3-9: Diagram of An EML Incorporating A DBR Laser and An EAM. (40298)

Mach-Zehnder Interferometers (MZIs) are used in a wide variety of applications within optics and optical communications. In the basic configuration, light is evenly split into two matched waveguides and then recombined at the interferometer output, as shown in Figure 3-10. If the phase in one of the waveguides is changed 180 degrees relative to the phase in the other, then destructive interference occurs at the output and the light intensity is at a minimum. The phase is changed through the electro-optic effect in which the RI of the material is changed with an applied electric field across one of the branches. When there is no phase difference, the light from the two branches adds constructively at the output and maximum light output is obtained. Mach-Zehnder Modulators (MZMs) are typically made with lithium niobate and can be operated at speeds beyond 100 Gbps, but again require 100s of volts to generate the required electric fields.

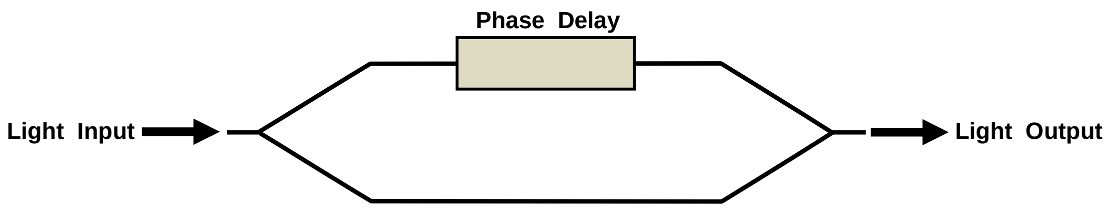

Figure 3-10: Diagram of An MZM. (40299)

3.4 Optical Amplifiers

An optical amplifier is a device that directly amplifies an optical signal without the need for O/E conversion. Most optical amplifiers are similar to lasers in that the amplification is achieved through stimulated emission within the device. The main difference between them is that optical amplifiers do not have mirrors forming an optical cavity. Optical amplifiers amplify all of the wavelengths in the gain spectrum of the gain medium (see Figure 3-1*(A)*), so they are particularly useful in WDM systems. Optical amplifiers can be built in three basic technologies; semiconductor laser technology, planar waveguide technology and optical fiber technology. Most use the principle of the laser in their operation (stimulated emission) but some use other principles, such as the Raman effect.

3.4.1 Erbium Doped Fiber Amplifier (EDFA)

Erbium doped fiber amplifiers are a specific case in the class of amplifiers known as rare earth doped fiber amplifiers (REDFAs). The most common dopants are erbium and praseodymium but other rare earth dopants, such as neodymium, can be used. Erbium atoms are added to the glass in a section of fiber. The fiber can be as long as 10 m and an EDFA can have up to 30 dB of gain. The EDFA requires a pump laser and also a coupler to combine the signal and the pump light as shown in Figure 3-11. A coupler at the output side can be used with a filter and detector to determine the output power level to provide feedback to the pump laser to control the gain and thus the output power level. The final piece in an EDFA is an isolator to prevent reflections from coupling back into the EDFA.

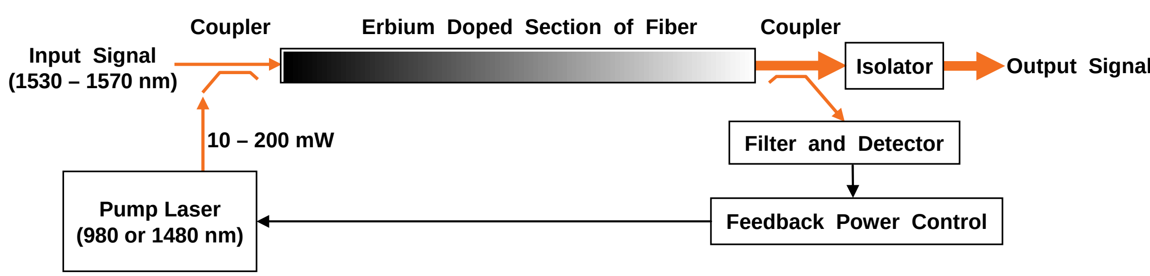

Figure 3-11: Schematic Diagram of An EDFA. (40540)

EDFAs use the same principle of operation as lasers. Electrons from the ground state are excited (pumped) to higher energy states in the erbium dopants in the fiber as shown in Figure 3-12. The lasers previously discussed in this reference guide are all electrically pumped. However, the EDFAs are optically pumped. So a light source, typically a laser, with a wavelength shorter than the signal to be amplified is used to excite electrons to higher energy states. Erbium has several unstable excited states and electrons can be excited to any of these states with the right wavelength laser, but a 980 nm laser is typically used. Electrons in these states decay to a metastable state without emitting photons. The electron energy given up in these transitions is transferred to the fiber in the form of phonons (lattice vibrations). From the metastable state, electrons either decay back to the ground state through spontaneous emission or through stimulated emission of photons. A population inversion (more electrons in the metastable state than in the ground state) is needed so that the light from the signal causes stimulated emission rather than being absorbed and exciting electrons from the ground state to the metastable state. The light from the stimulated emission has the same phase and direction as the light causing the emission, so the signal is amplified.

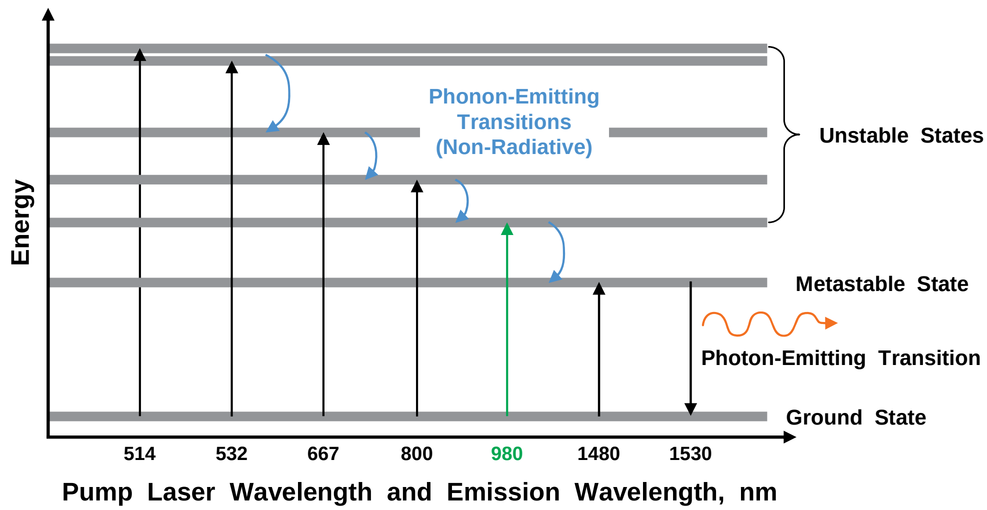

Figure 3-12: Energy States of An Erbium Doped Fiber. (40541)

Erbium doped amplifiers can also be fabricated in planar waveguide technology. These devices are referred to as erbium doped waveguide amplifiers (EDWAs). The majority of the literature on EDWAs is geared towards components for the metro and enterprise network applications. These devices have optical power gain in the range of 10 to 30 dB and waveguides lengths of approximately 7.5 to10 cm. The devices integrate the entire amplifier, including the pump laser and isolators, on a single chip. The next level of integration is to combine the EDWA with other components such as array waveguide gratings (AWGs). We have not been able to find any further information on EDWA development.

3.4.2 Semiconductor Optical Amplifier (SOA)

Semiconductor optical amplifiers (SOAs) are essentially FP lasers without an optical cavity. This is achieved by using low reflectivity mirrors, or no mirrors at all, at the ends of the device. Typically they are edge emitting devices with end facets made from the bare cleaved edges of the chip or with anti-reflections coatings on the cleaved edges. The gain can be increased by increasing the length of the device or by using higher reflectivity mirrors so the light passes through the gain medium several times before exiting the device. They can also be made with vertical cavity structures but the gain medium is so thin there is very little gain in a single pass through the device. So vertical cavity SOAs (VCSOAs) are made with higher reflectivity mirrors. In SOAs with mirrors on the ends, care must be taken to prevent the device from lasing.

SOAs are generally 400 µm – 2 mm in length with gain up to 30 dB for the longer devices. However, their output power is relatively low (several mW) and they have higher noise than EDFAs. As individual components, their power is further reduced by losses in the fiber pigtails. The main advantage of SOAs over EDFAs is that they are more easily integrated with other components in planar waveguide technologies. Figure 3-9 shows an example of an SOA integrated with a DBR laser and an EAM. For DWDM, an AWG with arrays of SOAs, lasers and modulators can be integrated on a single chip and an SOA and AWG with an array of photodiodes (PDs) can be integrated on a second chip, as shown in Figure 3-13, to fabricate a transmitter and receiver.

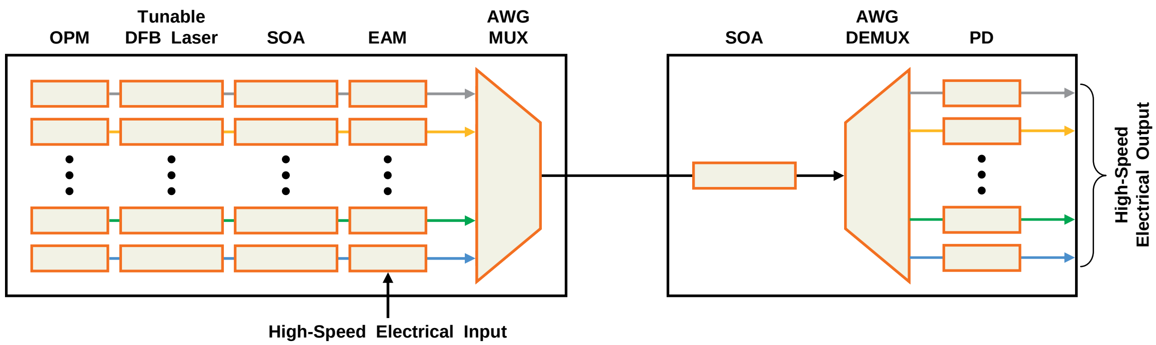

Figure 3-13: Diagrams of Possible DWDM Transmitter and Receiver Chips With Arrays of Devices.   (40542)

3.5 References

[3-1] Dutton, H. J. R.: Understanding Optical Communications, IBM Redbook 1998. http://www.redbooks.ibm.com/redbooks/pdfs/sg245230.pdf

[3-2] Li, H., and K. Iga (Editors): Vertical-Cavity Surface-Emitting Laser Devices. Springer Verlag, Berlin, 2002.

4 Connectors and Interconnection Hardware

There are many different styles of fiber optic connectors. Some are designed for a single fiber connection while others connect multiple fibers. Many connector styles can be specified for single-mode fiber (e.g., 6/125, 9/125 – core diameter in microns/cladding diameter in microns), or for multimode fiber (e.g., 50/125, 62.5/125). Certain connectors, such as the LC connector, are becoming more popular as termination densities increase. For example, the SC type is a snap-in connector that features a push-pull action versus a screw termination for ease of connection. The Electronic Industries Association (EIA) and the Telecommunications Industry Association (TIA) standard EIA/TIA 568B now allows any type of fiber optic connector design as long as it conforms to the Fiber Optic Connector Intermateability Standard (FOCIS) [4-1]. This is a component standard written for manufactures to ensure that their connector will mate with other manufacture's connectors of the same type. The following brief descriptions are of the most commonly used fiber optic connectors.

4.1 Standard Commercial Single-Way Connectors

FC

The FC/PC is a popular single mode connector. The connector is a screw-on mating style and is a robust and highly reliable design.

SC

The SC is a snap-in connector used in both single and multimode applications. It is a snap-in connector that latches with a simple push-pull motion. Through a snap clip it becomes a duplex version.

D4

The D4 connector, designed for durability, is made with a 2.5mm diameter ceramic ferrule. For repeatability and intermateability they have a high performance threaded mounting system and a keyed body.

ST

The ST is a bayonet mount with a long cylindrical ferrule to hold the fiber. The connector is spring-loaded, so you have to make sure it is properly seated. It is being replaced by the SC and LC connectors.

SMA

The SMA connectors are threaded connectors and used for connection of some medical equipment because of their low-cost multimode coupling. They are ideal for military applications. Unlike most connectors whose ferrules are ceramic or plastic, these ferrules are stainless alloy or stainless steel.

BICONIC

The BICONIC connector was developed in the late 1970s by Jack Cooper at Bell Labs. The connector has a cone-shaped ferrule that helps align the fiber at the connection interface. The robust connector design makes it well suited for military applications.

Images taken from : **http://www.fiberc.com**

4.2 Small Form Factor Connectors

The small form factor (SFF) connectors have been developed by various manufactures to achieve higher interconnect densities on distribution panels than would be possible using other standard connectors. The SFF are built around the 1.25mm ferrule instead of the 2.5mm ferrule used in FC, SC and ST types.

MT-RJ

The MT-RJ is a multimode molded duplex connector. It handles two fibers in its single ferrule. It is being replaced by the duplex LC connector.

LC

The LC is a new 1.25 mm ferrule connector that is half the size of an ST connector. It was introduced in the late 1990s to miniaturize connectors for higher density applications in patch panels or equipment. It has good performance and is the connector of choice for single-mode and high speed (Gb/s) networks.

MU The MU connector looks like a miniature SC with a 1.25 mm ferrule and is more popular in Japan.

E2000/LX-5 The E2000/LX-5 connector is like the LC connector but with a shutter over the end of the fiber.

Opti-Jack The Opti-Jack connector is a rugged duplex type that is designed around two ST-type ferrules in a package the size of a RJ-45 connector.

Volition The Volition connector is a simple and inexpensive duplex connector. The fibers are aligned in a V-groove like a splice, so no ferrule is used. The connector comes in both plug and jack versions but only the jack can be field terminated.

Images taken from: **http://www.fiberc.com**
and **http://www.thefoa.org/tech/ref/basic/term.html**

4.3 Blindmate Fiber Connector

Orthogonal Adapter Family Unique adapter assembly allowing for orthogonal (matrix) blindmate. Includes guidance and latching functionality

Image taken from: **http://www.amphenol-tcs.com**

4.4 Ribbon and Parallel Fiber Connectors

Fiber Gate

The Fiber Gate connector is an example of a parallel fiber connector. It will connect up to four optical fibers.

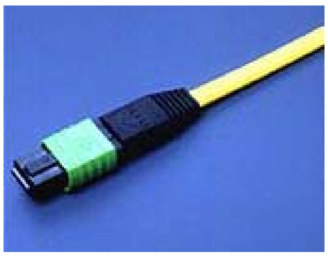

MPO/MTP

The MPO/MTP connector is an example of a ribbon fiber connector. It allows for a quick and reliable connection for up to 12 fibers. They are intended for installations that require many connections. However, because of the multiple alignment requirements, coupling loss is typically higher than it is for single-fiber connectors.

MPO Fan-Out

In the MPO Fan-Out device, one end of the ribbon cable is usually split to fan out each fiber with its own terminating connector.

Images taken from: **http://www.hubersuhner.com**
and **http://www.furukawa.co.jp/connector/mpo.htm**

4.5 Connector Ferrule Shapes and Polishes

The polish usually refers to the shape and finish of the ferrule (see Figure 4-1).

Early connectors were not keyed and the fiber ends could rotate with respect to each other causing abrasive damage in the mating adaptor. To prevent grinding scratches in the fiber ends an air gap was left between the connectors. Beginning with the SC, and later with FC connectors which were keyed to prevent rotation, the fibers were allowed to make physical contact (PC). Mating pressure was maintained with a spring.

The Air Gap connection type had an insertion loss of approximately 0.5 dB or more and a return loss of 20 dB. The PC type connector had an insertion loss typically of 0.3 dB and a return loss of 30 to 40 dB. By polishing the ferrule to a convex end for the Ultra Physical Contact (UPC) connector the insertion loss was reduced to less then 0.3 dB and the return loss improved to 40 dB or better. For single-mode systems that are extremely sensitive to reflections, the end of the ferrule is polished to 8 degrees in the Angled Physical Contact (APC) connector. Any reflected light at the interface is diverted into the cladding layer resulting in a return loss of 60 dB or better [4-2].

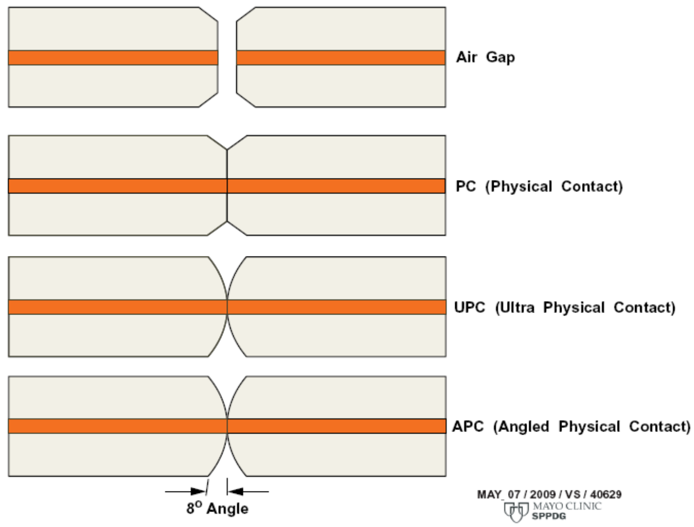

Figure 4-1: Connector Ferrule Shapes and Polishes (The Center Orange Strip Represents the Actual Optical Fiber Waveguide). (40629)

4.6 Abbreviation Meanings for Optical Connectors

Table 4-1 lists the abbreviations and their meanings for some of the more common optical connectors.

Abbreviation	Description or Meaning	Web Link
APC	Angled Physical Contact	
EIA	Electronic Industries Association	www.eia.org
FC	Ferrule Connector	
FOCIS	Fiber Optic Connector Intermateability Standard	
MPO/MTP	Multi-Fiber Push On	
MT-RJ	Mechanical Transfer Registered Jack	
PC	Physical Contact	
SC	Subscriber Connector, or Standard Connector	
SFF	Small Form Factor	
ST	Straight Tip	
TIA	Telecommunications Industry Association	www.tiaonline.org
UPC	Ultra Physical Contact	

Table 4-1: Meanings for Some of the More Common Optical Connector Abbreviations.

4.7 References

[4-1] VDV Works: Lennie Lightwave's New Guide to Fiber Optics. http://www.jimhayes.com/lennielw/llguide.pdf

[4-2] The Fiber Optics Association, Inc.: Reference Guide To Fiber Optics. http://www.thefoa.org

5 Analysis of Tradeoffs Between Optics and Copper

This section presents a framework for conducting a trade analysis of optical versus electrical high speed data communication links as applied to high performance compute (HPC) systems. The framework is based on key performance parameters that must be considered, including cost, which is probably the most difficult factor to predict quantitatively for systems of the future. A specific study comparing research implementations of optical and electrical interconnects for 10 and 20 Gbps data communications has recently been conducted at IBM [5-1] using similar performance metrics as the ones discussed below.

5.1 Performance Parameters

Bandwidth-Distance Product, B*d: This figure of merit refers to the useable distance, or reach, of an optical data communication channel taking into account the supported transmission bandwidth or data rate. Typically it is applied to multimode (MM) optical fiber, but the concept can be extended to other optical waveguide types and even to electrical links. Dimensionally it is specified in [MHz*km], or [GHz*km] when used in the context of longer links, but for shorter links where the relevant distances may be only 10's of centimeters, or less, the dimensions can be given in [GHz*cm].

When applied to a specific propagation medium, like MM optical fiber, the bandwidth-distance product, B*d, is a straightforward and well defined metric, specified by the product of the fiber length and the corresponding 3dB bandwidth of the optical signal (under specified launching and cabling conditions, at a specified wavelength). A practical rule of thumb is that the 3 dB bandwidth should be roughly 75% of the highest useable bit rate. Indeed, sometimes the "bit-rate-distance product" in [Gbps*m] is used instead. In MM fiber the signal degradation is primarily due to modal dispersion. Since the degree of modal dispersion is proportional to distance, one can assume the bandwidth-distance product to be a constant for a wide range of distances or bandwidths. It can be used, for example, to approximate the maximum distance that a specific fiber can support, given the bit-rate (or vice versa), and the B*d specification for the fiber. For example, for a 62.5-µm core MM fiber that is specified as having a B*d = 200 MHz*km at 850 nm, the maximum distance that this fiber would support a 1 Gbps data rate at 850 nm would be approximately (B*d/(bit-rate *0.75)) = (200/(1000*0.75)) = 0.267 km, or about 270 m.

However, the term, bandwidth-distance product is also used in other more loosely defined contexts including electrical links. For these cases the bandwidth-distance product is not necessarily a constant, i.e., the highest useable data rate may not be strictly inversely proportional to the reach of the link. In fact, there are important cases that deviate significantly from this linear relationship between bandwidth and distance. In SM fiber, for example, that does not employ any dispersion compensation, chromatic dispersion results in approximately a $B*d^{0.6}$ constant as indicated in Figure 5-1 [5-2]. For an electrical link (copper wire) where skin depth losses are dominant and no equalization circuitry is applied, the relevant constant is $B*d^2$, not B*d [5-3]. Nevertheless, for restricted ranges of either the bandwidth or distance (e.g., short distances relevant to HPC applications), the bandwidth-distance product is an approximate and useful figure of merit, for comparing different families of data communication systems including optical and electrical communication links.

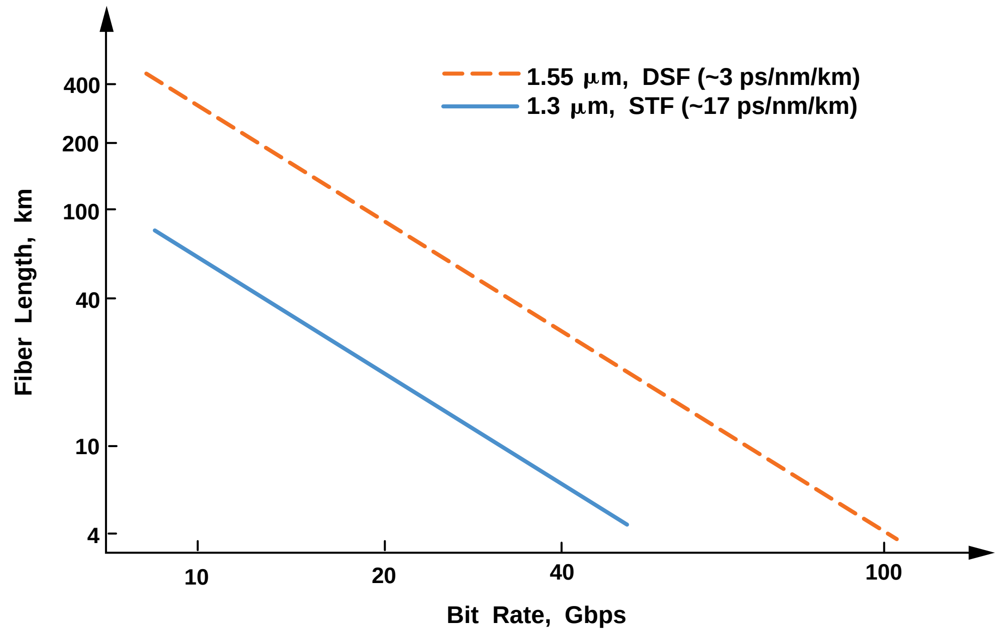

Reference: Grote, N., H. Venghaus (Editors): Fibre Optic Communication Devices, p. 152. Springer-Verlag, Berlin, 2001.

Figure 5-1: Bit Rate (Bandwidth)-Distance Limitation due to Chromatic Dispersion in Dispersion-Shifted Fiber (DSF) and Standard-Telecommunication Fiber (STF). (40884)

Bandwidth Linear Density, D: One of the primary advantages of optical waveguides which can be etched or printed on a printed circuit board (PCB) or silicon wafer, is their miniature size and inherently high signal isolation characteristics that allows dense spatial packing of multiple parallel waveguides. In some cases even orthogonal crossing of optical guides in the same layer can be realized with high isolation, something that cannot be achieved with electrical interconnects [5-4],[5-5]. Of particular interest is the use of densely packed optical guides for high aggregate bandwidth applications that would be used, for example, on a circuit board for relatively short reaches between ICs [5-6]. The metric of interest for this case is the total bit rate (bandwidth) linear density, with dimensions of [Gbps/mm]. An example that demonstrates the potential for high bandwidth density was shown in Section 2 (see inset in Figure 2-2) for embedded optical waveguides. For this case a 2D-array (4 layers of 12 guides each) of 48 optical waveguides are embedded in a distance of 2.75 mm on the circuit board. With each guide capable of supporting up to 12.5 Gbps, the bandwidth linear density is 218 Gbps/mm. Even higher linear densities can be achieved with more layers in the z-direction. An example of the higher channel density achievable in optics is shown in Figure 5-2 with regard to optical versus electrical connectors. Shown are two standard optical connectors (Xanoptix/Molex) that can carry at least 6 times the number of channels that the same cross-sectional area would be able to carry using a high speed electrical multi-channel connector (ERNI ERmet ZD).

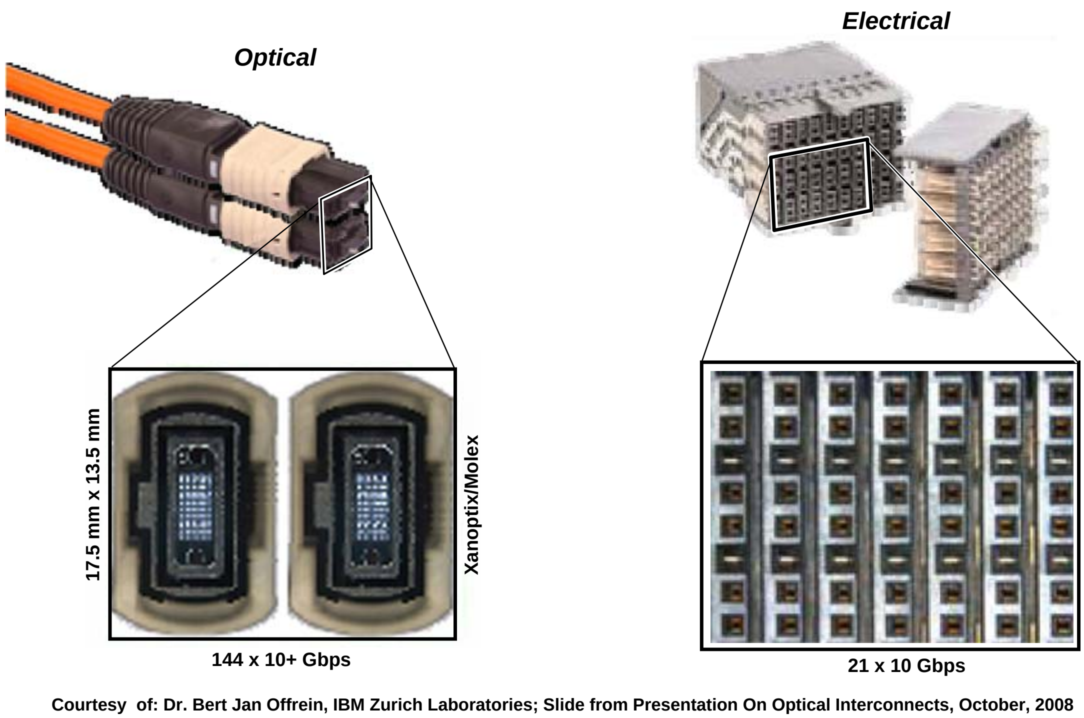

Figure 5-2: A Connector Comparison Showing the Increased Channel Density Possible With An Optical Connection Versus A High Speed Electrical Connection (More Than 6 x); the Optical Connector From Molex Incorporates A 72-Fiber (6 x 12) MT FerruleIn the Industry Standard MTP Connector; the ERNI ERmet ZD Electrical Connector Is Specifically Designed for Differential High-Speed Signal Transmission With Data Rates of Up to 10Gbits/sec. (41341)

Normalized Power Consumption, or Energy per Bit: A critically important consideration is the power consumption, or efficiency of a data communication link taking into account the bit rate. The relevant metric is measured in [mW/Gbps], or equivalently, in terms of the energy per bit, with dimensions [pJ/bit]. Examples of the power consumption for optical data communication links can be found in the data sheets for active optical cables (AOCs). AOCs incorporate an E/O and O/E transceiver at each end of the cable. As far as the user is concerned the optics is totally hidden and only electrical connections are required at the ends of the cable to electrically terminate the transceivers. The power consumed per "end" is equivalent to the power necessary to operate a unidirectional optical link, i.e., one transmitter and one receiver. The first four entries in Table 5-1 list representative vendors and basic power consumption characteristics for AOCs. Most of the cables employ vertical-cavity-surface emitting lasers (VCSELs) and MM fiber. However, Luxtera's cable employs edge-emitting lasers (EEL) and SM fiber. The last four entries in the Table are recent research results for a state-of-the-art optical and three all-electrical links. Although the electrical link at 6.25 Gbps has an attractive value for the power consumption metric, it should be noted that this experimental link operates only at a single custom data rate, it has smaller reach, and it lacks conformance to any communication standard. For more typical all-electrical SerDes-based chip-to-chip interconnects with reaches of less than 1m and with bit rates up to approximately 12 Gbps the normalized power consumption is usually greater than 10pJ/bit [5-7].

	Vendor or Reference	Wavelength [nm]	Channels x Bit Rate n x [Gbps]	Power per Transceiver [mW]	Energy/Bit [pJ/bit]	Approximate Reach [m]
Optical Links	Finisar	850 (VCSEL-based)	4 x 10 and lower rates	1500	38	100 on 4-lanes of MM fiber
	Emcore	850 (VCSEL-based)	4 x 5 and lower rates	1050	53	100 on 4-lanes of MM fiber
	Zarlink ("ZOE" SMT Component)	850 (VCSEL-based)	4 x 10 and lower rates	750	19	100 on 4-lanes of MM fiber
	Luxtera	1550 (DFB Laser-based)	4 x 10 and lower rates	2200	55	300 on 4-lanes of SM fiber
	Optical link research result (2007) from IBM (Terabus) 130 nm CMOS for drive electronics	985 (VCSEL-based)	16 x 10 (up to 16 x 15) and lower rates	2160 (at 10 Gbps)	13.5	≤ 1 On 16 lanes of embedded optical waveguides in FR4 (Optoboard)
Electrical Links	All-electrical research result (2007) from Rambus, Inc. in 90 nm CMOS	na	6.25 only	14	2.2	0.8 (200 μm microstrip on FR4 medium)
	All-electrical research result (2008) from Intel in 65 nm CMOS	na	1 x 15 and lower rates (Power Managed)	75	5	0.20 microstrip on FR4 medium
	All-electrical research result (2009) from Rambus, Inc. in 65 nm CMOS		512 x 16	172,000	21 (private communication)	0.08 microstrip on FR4 medium

DFB = Distributed Feedback
VCSEL = Vertical Cavity Surface Emitting Laser

Table 5-1: Power Consumption Characteristics for Active Cable Products and Recent Research Results for Optical and All-Electrical Transceivers (Data for Active Cables Are From Vendor Websites; Zarlink data is taken from ZOE evaluation board documentation; References For Last Four Rows Are [5-8],[5-9],[5-10], and [5-11], Respectively). (40342v4)

Reliability: The question of reliability of an optical versus an electrical data communication link is not so much a comparison exercise between technologies, as it is a determination of whether a particular implementation will have sufficient reliability for the expected life of the system in question. System life cycles may vary greatly ranging from 3 to 5 years for some commercial applications to 15 to 20 years for long-haul telecommunication systems. For HPC systems it is reasonable to assume a 10-year maximum component service life.

System reliability requirements dictate, to a large extent, the packaging and operating environment for the active devices, especially temperature, which in turn may have implications for overall module or transceiver power consumption. For example, wavelength stabilization and long term laser reliability in dense wave division multiplexing (DWDM) systems requires precise temperature control (often within tenths of a ^{o}C) and thermoelectric cooling (TEC) to lower laser operating temperatures and extend operating life. Both the TEC and the control electronics require extra power. In distributed feedback (DFB) lasers for telecommunications the power required for the TEC may be 60% (or higher) of the total laser transmitter power. In other words, reliability requirements may not only have implication on the type of device and packaging used, but also on overall power consumption.

The reliability of both VCSELs and edge emitting lasers have been studied and developed extensively over the last two decades (see Reference [5-12] for a review of VCSEL technology from a reliability perspective). Indeed, VCSELS as well as edge emitting lasers demonstrate lifetime characteristics similar to those observed in most semiconductor components. Namely, they follow three distinct rates of failure during their lifetime starting with: 1) After placing devices under stress (e.g., by operating them at elevated temperatures) one initially sees a high level of failures, which decreases fairly quickly over time. 2) There is then a period of a fairly low rate of failures, followed by 3) a significant increase at the point at which wear-out mechanisms dominate. Figure 5-3 depicts this type of "bathtub" curve. In an ideal world, one has a way of screening out the early defects ("infant mortality") such that they never reach the field, and the onset of wear-out occurs at a time longer than one expects the product to be in service. The very existence of such statistics for commercial laser devices is, to some degree, an indication of the technology's maturity.

VCSELs, for example, have been a commercial product since 1989, and in 2005, for example, it was estimated that over 50 million, 850 nm devices were shipped just in the first quarter of that year [5-13]. Although VCSEL reliability was an issue with earlier devices (primarily as a result of vulnerability to electrostatic discharge (ESD) and humidity), it is reasonable to say, that today, first-tier vendors have understood and adequately addressed these issues, (i.e., established the bounds of the bathtub curve) to where VCSELs can be applied to demanding, long life-cycle communication systems. Similarly, with edge emitters and DWDM systems, transceivers are being shipped today that are targeted for long haul telecommunications with required service lifetimes approaching 20 years [5-14].

For HPC applications, a specific study to estimate the reliability projections for densely-packed optical interconnects was conducted and reported in Reference [5-6]. In that study VCSEL-based interconnects were considered for reliability as a function of the number of I/O ports with and without several degrees of redundancy (see Figure 5-4). The argument is given that since the marginal cost of adding optical devices is low, then redundant lasers per channel can be incorporated to achieve higher system reliability. Based on a twenty-year average lifetime for a VCSEL, the curves indicate that a system with thousands of I/O could have lifetimes of 30 years or more.

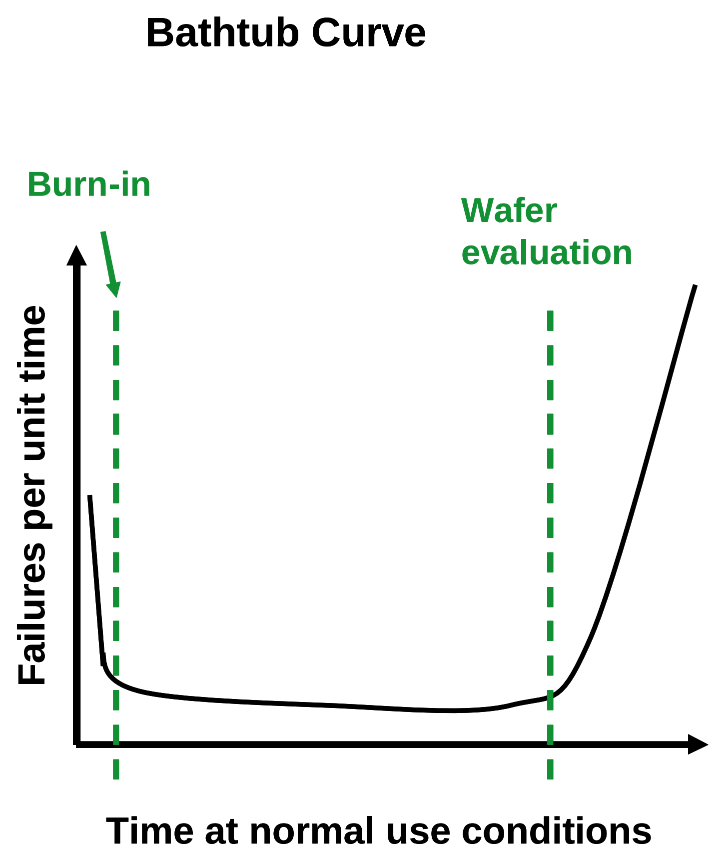

Figure 5-3: Illustration of the Reliability "Bathtub" Curve. (23546)

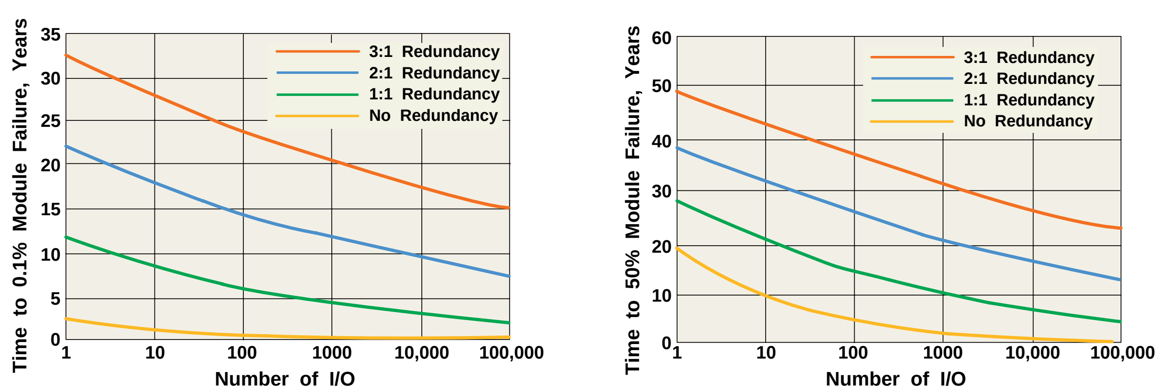

Courtesy of: Trezza, J., H. Hamster, J. Iamartino, H. Bagheri, and C. DeCusatis: Parallel Optical Interconnects for Enterprise Class Server Clusters: Needs and Technology Solutions. IEEE Optical Communications (February) 2003.

Figure 5-4: Projected Reliability for Dense VCSEL-Based Optical Interconnects (Based On 20-Year Medium Laser Lifetime). (40351v2)

Cost: It is well recognized in the telecommunications industry that because of the ever increasing demands for bandwidth (for services like online video, multimedia in general, wireless access and the availability of a wealth of devices that support such applications) the cost for telecommunication optical components and subsystems will of necessity have to continue to drop in order for the relevant service providers to stay economically viable. Likewise, there is a consensus among optical and photonic technology researchers, that the pressure to lower production costs will drive optical technology to miniaturization of optical components and on-chip integration, i.e., to the prolific use of photonic integrated circuits (PICs). It is estimated that the number of optical components (active and passive) that could be integrated in such a manner, could be in the several hundreds, if not thousands, per chip in the not so distant future [5-15].

It is reasonable to assume that PICs developed for telecommunications would also find their place in HPC systems. Both single-mode DWDM interconnect systems as well as parallel, densely-packed multimode interconnects have potential for intra-board data communication. In Reference [5-6], for example, the results of a cost study are presented with projected cost estimates for dense optical interconnects for future data communication systems. The estimates are given in normalized terms of cost per Gbps, and the cost per channel assuming a 10 Gbps data rate. Commodity-type volumes are assumed. The graph in that reference is reproduced in Figure 5-5 below. Clearly, the curves indicate the economy of scale effect and the reduction in optical interconnect costs with increased optical channel count. Miniaturization and the use of integrated photonics (nanophotonics) that allow the optical signals to be brought onto the PCB and routed right up to the electronic integrated circuit are the main factors for overall system cost reduction. Projected cost-per-channel for 10 Gbps channels and for number of channels on the order of 1000, is around 2½ cents. In contrast, today's transceiver module cost per channel, say for 10 Gbps Ethernet, is on the order of several dollars in high volumes [5-16]. Researchers from an HPC systems integrator company have stated that (as of four years ago) to be competitive with all-electrical copper channels the cost of optical links must come down below one dollar per Gbps [5-17].

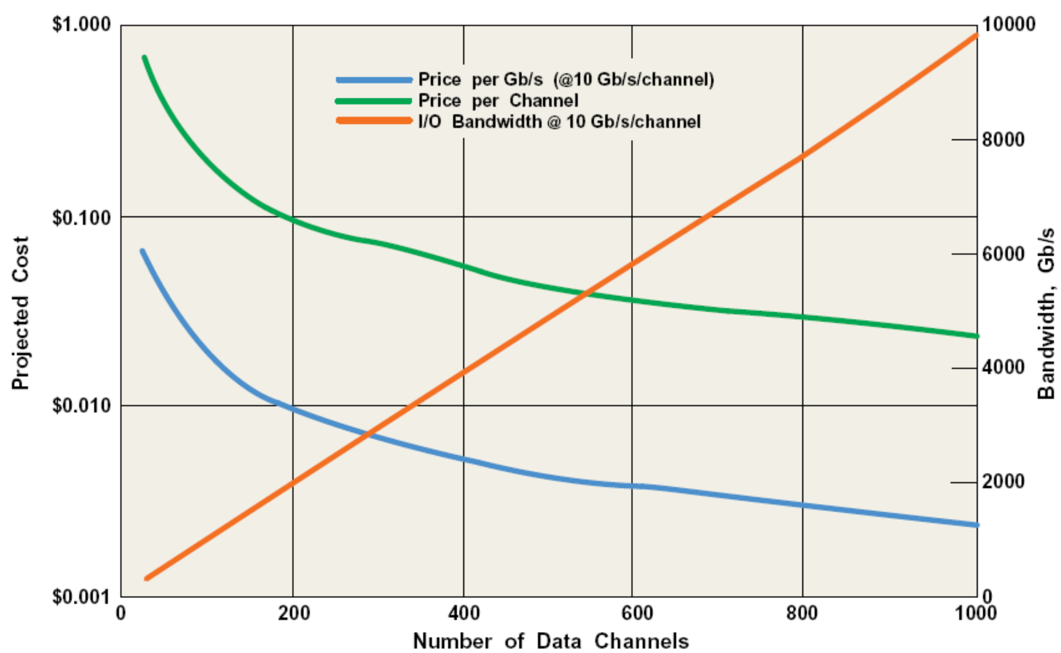

Figure 5-5: Projected Cost for I/O Bandwidth Possible in Large Volume using Dense Parallel Optics (Assumes On-PCB Optical Connections to the Electronic IC). (40339)

Performance Parameters That Are Especially Favorable for Optical Interconnects: Of the performance parameters discussed above, the bandwidth-distance product and the bandwidth density stand out as especially favorable for the case of optical interconnects for data communication above 10 Gbps and for intra-board interconnects, i.e., reaches of approximately several cm to 100 cm (and beyond if board-to-board or rack-to-rack communication are being considered). Assuming only skin depth losses and employing no equalization or error correction, an electrical channel's bandwidth limit degrades by a factor of 4 with every doubling of the distance [5-3], in contrast to the optical case where the bandwidth is nominally only halved. Furthermore, for a given bandwidth or bit rate (say 10 Gbps), the distance limit in multimode fiber, for example, might be hundreds of meters, while in the electrical domain it may be only several 10's of centimeters. i.e., the B*d constant itself is several orders of magnitude higher in optics. In addition, for an all-electrical channel that employs low cost media like FR4 there is also dielectric loss, which becomes significant at the higher frequencies and data rates so that the bandwidth in this case falls off even faster than when only skin depth loss is assumed [5-18].

From a practical, or bandwidth-distance perspective the "crossover zone" from electrical to optical links seems today to be in the vicinity of 400 to 800 Gbps-cm depending on the quality of the all-electrical channel media, starting at the low end with FR4 using no special equalization techniques. Applying equalization techniques the B*d product may be extended to around 1500 Gbps-cm, but at the expense of more power consumption and greater IC area for the electrical case [5-1]. For "single lane" bit rates that are much above 10 Gbps and for reaches greater than several ten's of cm, it is generally acknowledged that fiber-optic solutions will offer improved overall performance and a potentially lower cost solution than an all-electrical implementation if the electrical implementation is a standard printed circuit board material like FR4 [5-19]. A qualitative depiction of the crossover zone is shown in Figure 5-6 which is taken from the perspective of the printed circuit board industry [5-20]. On the optical side the comparison is made with respect to an optically-based PCB and assumes optical waveguides (glass or polymer) embedded in conventional circuit board material. As the figure shows, for commercial applications it will be the interconnect cost that will be the economic and dominant factor for deciding which technology is chosen for a given operating regime. This is also consistent with the opinion of the researcher cited in [5-17]. The figure also helps convey the notion that the electrical to optical transition is not an abrupt event, neither as viewed from the bandwidth-distance product nor from the time or cost perspective. In fact, the crossover zone is dynamic in that both electrical and optical technologies are continually advancing in performance and what was thought to be a clear optical application for an operating regime say, 10 years ago, may have an all-electrical solution today (e.g., application of sophisticated error correction and equalization approaches have pushed SerDes circuits operating near 20 GHz, albeit for short reaches).

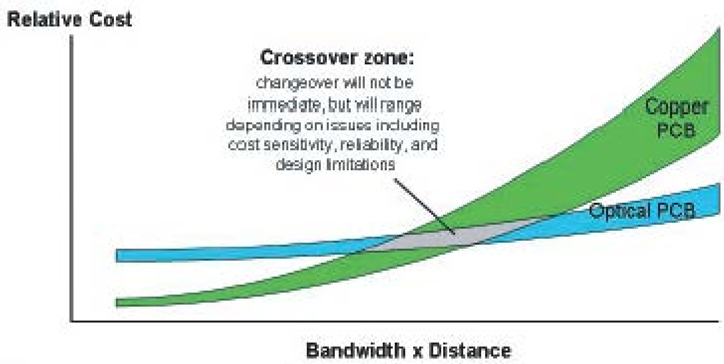

Figure 5-6: Optical Versus Electrical Technology Crossover As Applied To Printed Circuit Boards. (23538)

In terms of bandwidth density, optic solutions to intra-board parallel interconnects are especially attractive in that multiple high speed optical lanes embedded in a PCB (or flex material) can be tightly packed to achieve large aggregate bandwidths, on the order of TB/in or more. With the optics routed directly to the electronic integrated circuit, I/O density can be increased beyond that achievable with all-electrical I/O and pad ring limitations (especially when sophisticated error correction circuitry is employed, since it takes up more room on the IC).

5.2 Major Regimes of Operation for Optical Interconnects

It is clear today that in some applications optical communication links have an overwhelming advantage over their electrical counterparts, and in some cases have completely replaced the latter in actual fielded systems. An obvious example is the undersea long haul communication links which today use optical technology exclusively. At the other extreme it is likewise evident that on-chip high bandwidth communications are today implemented with electrical interconnects. In between there are other important operating regimes to consider including those applicable to HPC systems that are, in some cases, just beginning to apply optical technology. Table 5-2 (courtesy of Dr. Bert Offrein, IBM Zurich) categorizes these different operating regimes according to approximate link distances. In general, optical technology is replacing high speed electrical links for shorter and shorter distances. Today rack-to-rack communications for large computer systems, where the distances are on the order of 30 meters or less, are largely implemented with optical links (see next subsection). However, board-to-board, and especially on-board links with distances of approximately 1m or less are operating regimes that, although typically are today implemented with electrical links, are being seriously “challenged” by optical interconnect options. In this sense this intra-board and inter-board regime is today in the midst of the crossover zone between optical and electrical technologies. In terms of the key performance parameters, Table 5-3 gives an approximate comparison of electrical versus optical interconnects for the intra-board and inter-board operating regimes for data rates around 10 Gbps per channel based on specific research implementations (using multimode polymer waveguides for the optical case) as reported in the literature.

Finally, it may be noted that some researchers have doubts whether optical links for intra-chip communications will ever become economically viable.

	Internet, Wide Area Network	Local Area Network	Rack-to-Rack	Card-to-Card	On-Card	On-MCM	On-Chip
	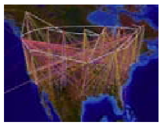				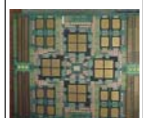	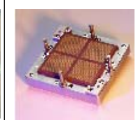	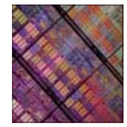
Distance	multi-km	10-2000 m	30+ m	1 m	0.1-0.3 m	5-100 m	0.1-10 m
Number of Lines	1	1-10	~100	~100-1000	~1000	~10,000	~100,000
Use of Optics	Since the early 80s and the early 90s	Since the late 90s	Now	2010+	2010-2015	Probably After 2015	Later, if ever

Courtesy of: Dr. Bert Jan Offrein, IBM Zurich Laboratories; Slide from Presentation On Optical Interconnects, October, 2008

Table 5-2: Regimes of Operation for Optical Interconnects in Terms of Reach and Number of Interconnects. (40369)

Performance Parameter	Electrical	Optical	References/Notes
BW*Distance [GHz*cm]	~ 350 (on FR4 – no EQ or EDC); Up to 1200 with EQ and EDC	> 4500 (multimode system)	[5-1],[5-19]
Linear BW Density [Gbps/mm]	25	192 (multimode system)	[5-1],[5-19]
Normalized Power Consumption [pJ/bit]	~ 5– 20 (best numbers)	~ 13 – 60 multimode system)	[5-1],[5-8],[5-10] At times difficult to assess how much of the circuitry is being accounted for
Reliability	Sufficient for 10 - 20 yr lifetimes	Comparable to electrical	Laser reliability not an issue as it was 10 years ago

Performance Parameter	Electrical	Optical	References/Notes
Cost	Several dollars per transceiver (today)	Depends on future HPC market and volume demand	Potential for less than $0.50 per transceiver in large volumes[5-6],[5-19]

EQ: Equalization
EDC: Error Detection and Correction

Table 5-3: Comparison of Key Performance Parameters for Optical and Electrical High Speed (> 10 Gbps) Data Communication Links Relevant to On-Board and Between-Board Interconnects (For Approximate Reaches of Several cm's to 1 Meter).

5.3 A Nearly Complete Electrical-to-Optical Crossover Example: High Speed Communication Links in Enterprise Data Centers

Twenty years ago, most data centers used electrical cabling for their infrastructure, for both communications and cabinet interconnections. Optical connections were primarily employed for communications outside the data center, due to their distance, or reach advantage over copper interconnects (which was limited to approximately 100 meters).

Optical interconnections started appearing within data centers for high-speed Storage-Area Networks (SANs), utilizing Fibre Channel (FC) as the protocol for connecting centralized storage to servers, and for interconnecting storage disk drives/shelves and their servers. The primary driver for using optical fiber as the technology of choice in the data center, in the early days, was bandwidth – higher bandwidths on optical media versus electrical interconnects allowed data to flow more efficiently to and between servers. Today optical and copper both support increasingly higher speeds (active optical cables with aggregate bandwidths of up to 40 Gbps are now becoming commercially available) to the point where speed isn't necessarily the driving factor for choosing optical technology over copper.

Why are data centers continuing to bring in optical connections? Optical connections allow the use of much smaller connectors, utilizing less backplane real estate. This increased connector density allows for more connections within a cabinet, which leads to more cable runs both intra-cabinet and inter-cabinet. Optical cables are a physically smaller cable, which becomes increasingly important when one starts utilizing hundreds of cables. Smaller cable bundles also tend to block airflow less, leading to more efficient cooling. Optical cables are not subject to electromagnetic crosstalk (due to close proximity to power cables or other signal cables) like copper cables can experience; this becomes more of a concern with all electrical interconnects as speeds increase.

Data centers are also becoming more connected to other data centers, for increased redundancy and disaster recovery. These connections are typically over optical fiber for speed and distance (> 100M) reasons, and bringing optical connection all the way down to the server (instead of to a demarcation point that converted to copper) is becoming increasingly common.

5.4 Economic and Other Factors to Consider When Assessing "Optical Versus Copper"

The trade off considerations discussed above have focused primarily on technological factors. However, other less-technical considerations may be as important, if not critical to the assessment. In this context it is reasonable to ask the question, given the state-of-the-art in optical and photonic technology, why has more of it not been adopted in short link regimes like the typical board-to-board and back-plane configurations for HPC systems? Even a simplified answer is multi-faceted and involves technological, business, and simply practical considerations. First of all, it should be emphasized that optics has been adopted in rack to rack interconnects using commercially available active optical cables as the previous subsection has indicated. One reason for their success is that as far as the end user is concerned, the active cable optics is completely isolated or hidden within the cable; the interfaces at both ends being strictly electrical. However, in many cases the optical technology has tended to stay in the laboratory and not migrate into actual HPC system upgrades. To some extent, electrical solutions have a better chance for acceptance since they have the advantage of electrically-based legacy systems that are easier to modify with advanced electrical solutions rather than updating them with an altogether different technology. Also, the market crash in late 2000, tended to suppress a lot of the photonics research efforts that were prevalent at that time, preventing them from reaching fruition and widespread system insertion. Today's economy (2009, 2010) is likewise unfavorable. Another practical, but often critical missing ingredient, has been compliance of the newer optical components and subsystems to established telecommunication industry standards. Indeed, some standards have not yet even been defined. For example, the IEEE group is still working on the standard for 40-Gbps Ethernet. Nevertheless, it is expected to have final approval for the technology by the international standards body by early 2010.

Other barriers to the introduction of optics are economic or market related. The cost of optical links is still too high to be economically viable for widespread use in HPC systems. A large volume commercial market has not yet emerged for HPC application and thus the economics of scale cannot be taken advantage of. In this sense the industry faces a "chicken and the egg" dilemma. For example, while IBM's optoboard has been demonstrated in an experimental setting, there is no current manufacturing infrastructure to support it and make it commercially available to potential system and subsystem developers [5-21]. Consequently, there is no high volume driver to lower the cost. In fact, there are other competing approaches that are currently being developed in various research laboratories. All these factors have contributed in their own way to delaying introduction of photonic and optical technology at least at the board-to-board and chip-to-board levels.

5.5 References

[5-1] Pepeljugoski, P., M. Ritter, J. A. Kash, F. Doany, C. Schow, Y. Kwark, L. Shan, D. Kam, X. Gu, and C. Baks: Comparison of Bandwidth Limits for On-card Electrical and Optical Interconnects for 100 Gb/s and Beyond. Proceedings of the SPIE, Volume 6897 (January) 2008.Grote, N., and H. Venghaus (Editors): Fibre Optic Communication Devices. Springer-Verlag, Berlin, 2001.

[5-3] Dally, W. J., M-J. E. Lee, F-T. An, J. Poulton, and S. Tell: High Performance Electrical Signaling. Massively Parallel Processing, 1998. Proceedings. Fifth International Conference, pp. 11-16 (June) 1998.

[5-4] Poon, A. W., et al.: Cascaded Microresonator-Based Matrix Switch for Silicon On-Chip Optical Interconnection. Proceedings of the IEEE, 97(7):1216-1238 (July) 2009.

[5-5] Bamiedakis, J., et al.: Low Loss and Low Crosstalk Multimode Polymer Waveguide Crossings for High-Speed Optical Interconnects. Conference on Lasers and Electro-Optics, 2007. CLEO 2007, pp.1-2 (May) 2007.

[5-6] Trezza, J., H. Hamster, J. Iamartino, H. Bagheri, and C. DeCusatis: Parallel Optical Interconnects for Enterprise Class Server Clusters: Needs and Technology Solutions. IEEE Optical Communications, pp. s36–s42 (February) 2003.

[5-7] Alduino, A.: Implementation Challenges of Future Optical Interconnects. Presentation at the SMTA International (October) 2007.

[5-8] Doany, F. E., et al.: Terabus: A 160-Gb/s Bidirectional Board-Level Optical Data Bus. Conference Proceedings of the 20th Annual Meeting of the IEEE Lasers and Electro-Optics Society, pp. 545-546 (October) 2007.

[5-9] Poulton, J., R. Palmer, A.M. Fuller, T. Greer, J. Eyles, W .J. Dally, and M. Horowitz: A 14 mW 6.25-Gb/s Transceiver in 90-nm CMOS. IEEE Journal of Solid State Circuits, 42(12):2745-2757 (December) 2007.

[5-10] Balamurugan, G., et al.: A Scalable 5-15 Gbps, 14-75 mW Low-Power I/O Transceiver in 65 nm CMOS. IEEE Journal of Solid State Circuits, 43(4):1010-1019 (April) 2008.

[5-11] Beyene, W. T., et al.: The Design and Signal Integrity Analysis of a TB/sec Memory System. DesignCon 2009 Conference (February) 2009.

[5-12] Mayo Final Report: A Study of Optical versus Electrical Interconnects for HPCS, Report Number Mayo-R-08-16-R0 (May) 2008.

[5-13] WaveReport, VCSEL Tutorial: http://www.wave-report.com/tutorials/vcsel.htm

[5-14] Private Communication, Infinera Technical Representative, 2009.

[5-15] Smit, M.: Comment made at the 2008 OIDA Conference in Monterey, CA.

[5-16] Patel, A.: Gigabit Ethernet Transceivers Ignite Port-Count Explosion. CommsDesign, August 10, 2005. http://www.commsdesign.com/design_corner/showArticle.jhtml?articleID=167600702

[5-17] Private Communication, G. Astfalk, Chief Scientist, Hewlett Packard (HP).

[5-18] Yuceturk, E., and S. Esener: Comparative Study of Very Short Distance Electrical and Optical Interconnects Based On Channel Characteristics. Optics in Computing (Trends in Optics and Photonics Series Vol. 90), Optical Society of America. 2003, pp.7-9. Washington, DC, USA (June) 2003.

[5-19] Pepeljugoski, P., F. Doany, D. Kuchta, L. Schares, C. Schow, M. Ritter, and J. Kash: Data Center and High Performance Computing Interconnects for 100 Gb/s and Beyond. OFC 2007 paper OMR4, 2007.

[5-20] Fisher, J.: Optoelectronic Substrates: Will They Happen. Journal of Circuits Assembly (November) 2007.

[5-21] Private Communication: Doany, F., and J. Kash, IBM, Watson Research Center (April) 2008.

6 Existing Optical Technology

There is an impressive array of optical products currently available, much of which is not relevant for high performance compute (HPC) applications. Many of these are available in a variety of packages and configurations as well. This section is not meant to be an exhaustive catalog of all available optical components. Rather it is a short listing of components to give a representative example of what is currently available. The information was collected through web searches based on component names and by searching vendor websites for data sheets.

6.1 Optic Fiber and Waveguides

6.1.1 Single-Mode Fiber (SMF)

Single-mode fiber (SMF) can be used over a range of wavelengths in the near-infrared region of the spectrum. It is typically used in long-haul communications at wavelengths around 1550 nm because of the extremely low attenuation at these wavelengths. Table 6-1 lists typical characteristics of several single-mode fibers. The cutoff wavelength is the minimum wavelength at which the fiber will support only one propagating mode. Light at wavelengths below the cutoff wavelength may experience multimode operation and additional sources of dispersion. The mode field diameter is the diameter of the area which carries the optical signal in a single-mode fiber. The majority of the light is carried in the core of the fiber but some of it is carried in the cladding.

Vendor	Operating λ (nm)	Cutoff λ (nm)	Mode Field Diameter (μm)	Attenuation (dB/km)	Numerical Aperture
Corning	1260-1625	≤ 1260	9.2 ± 0.5 @ 1310 nm 10.7 ± 0.5 @ 1550 nm	0.35 – 0.40 0.17 – 0.18	-
Corning*	1260-1625	≤ 1260	9.2 ± 0.4 @ 1310 nm 10.4 ± 0.5 @ 1550 nm	0.33 – 0.35 0.19 – 0.20	0.14
Fibercore	830	730 ± 70	5.6	< 5	0.12
Nufern	780-970	730 ± 30	5.0 ± 0.5 @ 850 nm	< 3.5 @ 859 nm	0.13
Nufern	1060-1600	920 ± 30	6.2 ± 0.5 @ 1060 nm	< 1.5	0.14
Nufern	1300-1625	1260 ± 30	8.6 ± 0.5 @ 1310 nm 9.7 ± 0.5 @ 1550 nm	≤ 0.5 ≤ 0.5	0.13
Nufern	1460-1620	1400 ±50	9.5 ± 0.5 @ 1550 nm	< 0.5	0.13

Table 6-1: Single-Mode Fiber Characteristics. Core Diameter ~ 8.2 μm, Cladding Diameter = 125 μm, Coating Diameter = 245 μm. * Industry Standard.

6.1.2 Multimode Fiber (MMF)

Multimode fiber (MMF) is typically used with 850 nm wavelength light in shorter distance applications (< 1 km). This is due to its higher attenuation which results from higher dopant levels in the fiber. The advantage of MMF is that it is easier to align it to other optical components because of its larger core diameter. Table 6-2 and Table 6-3 show characteristics of several graded-index MM fibers and step-index MM fibers, respectively.

Vendor	Core Diameter (μm)	Bandwidth-Distance Product (MHz-km)	Attenuation (dB/km)	NA
Corning	50.0	4700 @ 850 nm	≤ 2.3 @ 850 nm ≤ 0.6 @ 1300 nm	0.20
Corning	50.0	2000 @ 850 nm		0.20
Corning	50.0	850 @ 850 nm		0.20
Corning	50.0	510 @ 850 nm		0.20
Corning	62.5	385 @ 850 nm	≤ 2.9 @ 850 nm ≤ 0.6 @ 1300 nm	0.275
Corning	62.5	220 @ 850 nm		0.275
Thorlabs	62.5	160 – 400 @ 850 nm 300 – 1200 @ 1300 nm	2.7 – 3.2 @ 850 nm 0.6 – 0.9 @ 1300 nm	0.275

Table 6-2: Graded-Index Multimode Fiber Characteristics. Cladding Diameter = 125 μm, Coating Diameter = 240-250 μm.

Vendor	Core Diameter (μm)	Operating λ (nm)	Minimum Bend Radius (mm)	Numerical Aperture
Nufern	105	800 – 1600	25	0.12
Nufern	105	800 – 1600	25	0.15
Nufern	105	800 – 1600	25	0.22
Thorlabs	105	400 – 2400	21	0.22
Thorlabs	50	400 – 2400	10	0.22

Table 6-3: Step-Index Multimode Fiber Characteristics. Cladding Diameter = 125 μm, Coating Diameter = 240-250 μm.

6.1.3 Plastic Optical Fiber (POF)

Plastic optical fibers (POFs) are an alternative to glass-based fibers. They typically have higher attenuation than glass fiber but are generally easier to work with and are lower cost. Table 6-4 shows typical characteristics of several POFs. They are also available from several other vendors but these fibers generally have higher attenuation than those listed.

Vendor	Core Diameter (µm)	Bandwidth-Distance Product (MHz-km)	Attenuation (dB/km)	Minimum Bend Radius (mm)	Numerical Aperture
Chromis	50	> 300 @ 850 nm	< 60 @ 850 nm < 60 @ 1300 nm	5	0.190
Chromis	62.5			5	0.190
Chromis	120			10	0.185
FiberFin	750	150 @ 650 nm	< 200 @ 650 nm	25	-
Fuji Film	120	500 @ 850 nm	100 @ 850 nm	10	-
Mitsubishi	486	-	< 210 @ 650 nm	-	0.3
Mitsubishi	980	10 @ 650 nm	< 160 @ 650 nm	25	0.3

Table 6-4: Plastic Optical Fiber Characteristics.

6.1.4 Specialty Fiber

A variety of specialty fibers are available from many different vendors. These include a variety of polarization maintaining fibers, (chromatic) dispersion compensating fiber and other types. In the interest of space, specialty fibers will not be discussed here.

6.1.5 Planar Waveguide Technologies

There are many universities and companies working on optical waveguides (non-fiber) on a variety of substrates. These are generally targeted for on-board or on-chip applications. Luxtera and Infinera have developed waveguide technology to interconnect on-chip devices for photonic lightwave circuits (PLCs) on silicon and photonic integrated circuits (PICs) on InP, respectively. The only difference between PLCs and PICs is that PLCs are on silicon and PICs are on other semiconductor substrates. Generally for these waveguides to be useful, the technologies must also support light sources, modulators and detectors. Other companies are working on polymer based waveguides laminated onto FR4 and other substrates for on-board optical interconnects. PPC Electronic AG is developing glass waveguides with polymer cladding on a glass substrate. Table 6-5 summarizes some of the characteristics of these technologies.

Vendor	Technology	Typical Propagation Loss (dB/cm)	Core Width / Waveguide Size ($\mu m/\mu m^2$)	Bottom Oxide/ Substrate (nm/-)	Type
[6-1]	Silicon	0.1 – 3.0	-	-	-
[6-2], [6-3], [6-4]	Silicon MM	0.1 – 0.9	-	-	-
Luxtera	Silicon MM	0.1 – 0.9	3.0	800	Ridge
Luxtera	Silicon SM	1.0 – 1.7	0.5	800	Ridge
IBM	Silicon SM	3.6	0.445	2000	Strip
University of Lyon	Amorphous Silicon SM	4-5	< 0.5	-	-
[6-1], [6-5]	Silicon Photonic Crystal	20-30	-	-	-
Refs in [6-6]	InP	0.2 – 2.0	-	-	-
Infinera	InP	0.6 – 0.7	-	-	-
PPC Electronic	Glass	0.05-0.20	80	Glass	Embedded core
Information and Communications University	Polymer	0.25	100 x 60	270 um Silica glass plate	-
IBM	Polymer	0.04-0.05	35 x 35	FR4	-

Table 6-5: Waveguide Characteristics.

6.1.6 Fiber Optic Appliqués

In addition to waveguides on rigid substrates such as silicon, glass and FR4, waveguides have also been implemented on flexible substrates, such as kapton. These components are commonly referred to as appliqués. The waveguides are generally either standard optical fibers attached to the substrate or polymer waveguides fabricated on the substrate. The ends of the waveguides are typically connectorized for easy connection to other components. Figure 6-1 shows two examples of appliqués with optical fiber and polymer waveguides.

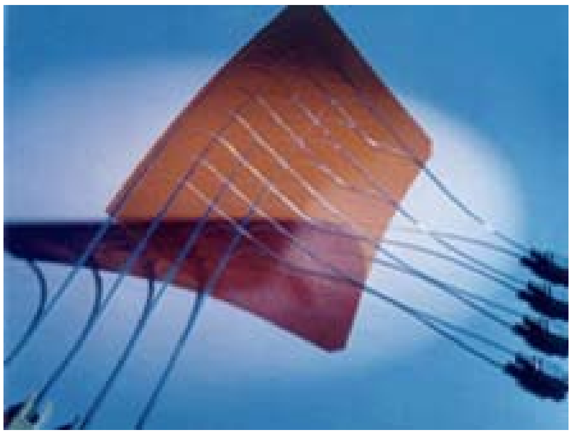

Molex, FlexPlane™ Optical Circuit, 8 x 8 Perfect Shuffle, 74 mm Wide, 137 mm Long, 1.5 mm High

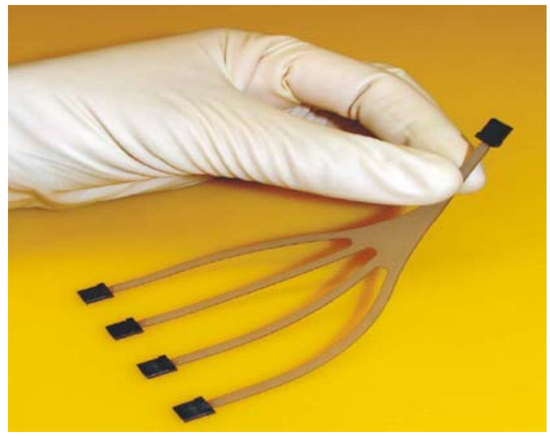

IBM-Zurich, Waveguide Flex, 4 x (1x12) (250 μm pitch) and 1 x (1x48) (62.5 μm pitch)

Figure 6-1: Two Examples of Optical Waveguides on Flexible Substrates. (40927)

6.2 Light Sources

Laser diodes are available from many different vendors at many different wavelengths. They are also available as bare die, in packages or mounted in subassemblies with drivers. Figure 6-2 shows examples of laser diodes mounted in packages and transmitter subassemblies. There are too many options to list all of them here so only a few packaged lasers will be highlighted at wavelengths of 850 nm (Table 6-6) and 1550 nm (Table 6-7). These lasers can also be obtained in transmitter subassemblies with thermoelectric coolers (TECs), modulators, drivers and additional control electronics. The information here was obtained from data sheets that were found either at the company website or at http://www.alldatasheet.com. Some of the laser diode vendors are Agere Systems, Bookham, Eagleyard Photonics, Eudyna, Fujitsu, Furukawa, Hitachi, JDS Uniphase, Laser 2000, Mitsubishi, NEC, NTT/NEL, Optoway, Pirelli, Rohm, Santur, Sanyo, Sharp and Thorlabs. While there are numerous papers regarding the development of DBR lasers and several commercial DBR lasers at shorter wavelengths (< 1100 nm; http://www.eagleyard.com/en/products/dfb-dbr-laser/), we could not find any commercial DBR lasers at 1550 nm.

Thorlabs

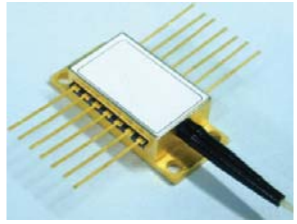

Thorlabs

LEFT: Laser in 5.6 mm package (typically with photodiode to monitor laser output)

RIGHT: Laser in pig-tailed butterfly package (typically with photodiode to monitor laser output and with thermo-electric cooler)

Multiplex

Bookham

Two examples of tunable transmitter arrays

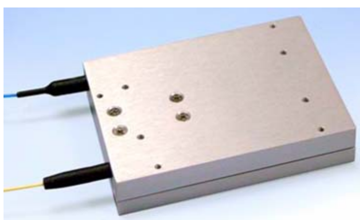

Avanex

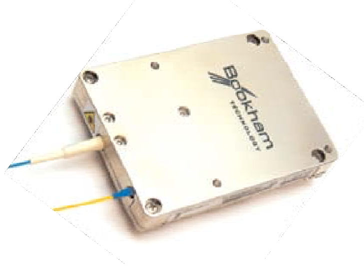

Bookham

Two examples of tunable small form factor transponders

Pictures Are Not To Scale

Figure 6-2: Examples of Currently Available Packaged Laser Diodes and Transmitter Assemblies. (40920)

Vendor	Type	Threshold Current (mA)	Maximum Optical Power (mW)	Typical Electrical Power (mW)	Beam Divergence (Degrees)
Bookham	Single-mode VCSEL	2.5	1	< 16	17
Bookham	Multimode VCSEL	1.0	6	12	32
Thorlabs	Multimode VCSEL	2.2	> 2	15	25
Thorlabs	Single-mode FP	25	10	45	30

Table 6-6: 850 nm Laser Diode Characteristics.

Vendor	Type	Threshold Current (mA)	Maximum Optical Power (mW)	Typical Electrical Power (mW)	Beam Divergence (°)
Laser 2000	FP	12	10	48	40
Mitsubishi	FP	10	6	33	30
Mitsubishi	DFB (InGaAsP)	8	10	28	35
NEC	DFB (InGaAsP)	12	10	-	30
Agere	DFB	35	30	625	-
Optoway	DFB	10	6	36	35

Table 6-7: 1550 nm Laser Diode Characteristics.

6.3 Modulators

The two main types of modulators being used in telecommunications are the lithium niobate Mach-Zehnder interferometer modulator and the semiconductor electro-absorption modulator (EAM). EAMs are typically integrated with distributed feedback (DFB) lasers to reduce loss and assembly cost, so we could not find any discrete EAMs. Table 6-8 gives the characteristics of several lithium niobate Mach-Zehnder interferometer modulators which operate up to 12.5 Gbps.

Vendor	Extinction Ratio (dB)	Insertion Loss (dB)	RF Drive Voltage (V)	$V\pi$ Bias Voltage (V)
Avanex	≥ 20	3.5	-	5
Avanex	≥ 20	4	-	2
Avanex	≥ 13	≤ 5	-	-
JDSU	25	2.5 – 5	2.6	2
JDSU	27	4.5	5.5	4
JDSU	27	4.5	4.5	2.5

Table 6-8: Modulator Characteristics. V_π Is The Half-Wave Voltage Which Is The Voltage Required To Switch From One State To The Other (e.g. From On To Off).

6.4 Active Optical Cables

Active optical cables incorporate an E/O and O/E transceiver at each end of the cable, so the inputs and outputs of the cables are electrical with a fiber optic cable in the middle. There are several active optical cables available on the market. They are typically designed to be used for a particular protocol (e.g. Infiniband) for short communications links (< 200 m). Table 6-9 lists the characteristics of several active optical cables.

Vendor	Channels x Speed (- x Gbps)	Aggregate Speed (Gbps)	Power (W)	Energy/ Bit (pJ/bit)	Fiber (SMF / MMF)	Form Factor	Reach (m)
Emcore	(4+4) x 3.2	25.5	0.8	31.4	MMF	MDI	300
Emcore	(4+4) x 5	40	1.0	25.0	MMF	MDI	150
Finisar	(1+1) x 10.3	20.6	1.0	48.5	MMF	?	30
Finisar	(4+4) x 10	80	3.0	37.5	MMF	QSFP	100
Finisar	(12+12) x 12.5	300	6.0	20.0	MMF	CXP	100
Fujitsu	(4+4) x 3.2	25.5	0.9	35.3	MMF	MDI	300
Luxtera	(4+4) x 10.3	82.4	?	?	SMF	QSFP	4,000
Reflex Photonics	(4+4) x 5	40	2.6	65.0	MMF	QSFP	200
Reflex Photonics	(4+4) x 10	80	2.6	32.5	MMF	QSFP	100
Tyco	(4+4) x 5	40	2.6	65.0	MMF	QSFP	100
Tyco	(4+4) x 10	80	3.0	37.5	MMF	QSFP	100
Zarlink	(4+4) x 5	40	< 2	< 50	?	QSFP	60
Zarlink	(4+4) x 5	40	2	50	?	CX4	100

Table 6-9: Active Optical Cable Characteristics.

6.5 Optical Amplifiers

The three main types of optical amplifiers are semiconductor optical amplifiers (SOAs) and erbium-doped fiber and waveguide amplifiers (EDFAs and EDWAs). Several characteristics of these optical amplifiers are listed in Table 6-10. This information was obtained from the company websites. The majority of these amplifiers are designed for the telecom industry so they are relatively large (rack mounted) and their power consumption is high (> 5 W).

Vendor	Type	Gain (dB)	Input Power (dBm)	Output Power (dBm)	Power (W)	Noise Figure (dB)	Volume (mm^3)
Avanex	EDFA	15 – 30	-20 – 0	9 – 15	1.2	5.5 – 7.0	40 x 70 x 7.5
Avanex	EDFA	23	-25 – -1	20.5	9.0	5.5	70 x 90 x 15
Avanex	EDFA	24 – 34	-36 – -3.5	20.5	14	5.0 – 7.5	212 x 130 x 20
QAMnet	EDFA	< 45	-15 – +5	15 – 24	60	5.2	483 x 356 x 44
QAMnet	EDFA	-	-10 – +8	15 – 24	35	5.0	254 x 152 x 33
Red-C	EDFA	8 – 35	-31 – +4	22	20	5.5	130 x 200 x 25
Tuolima	EDFA	-	-3 – +10	14 – 24	-	5.0	483 x 381 x 45
Inplane Photonics	EDWA	20	-	-	-	6.0	50 x 10 x 6
Alphion	SOA	19	-	8	4	7.5	30 x 12.7 x 8.4
Alphion	SOA	10	-	12	4	-	30 x 12.7 x 8.4
CIP	SOA	22	-15 – +13	8	< 7.6	7	30 x 12.7 x 8.4
CIP	SOA	14	? – +13	16	< 9.0	6.5	30 x 12.7 x 8.4
CIP	SOA	30	-10 – +13	5	< 7.6	10	30 x 12.7 x 8.4
CIP	SOA	30	? – +13	5	< 0.2	6	5.6 mm dia.
Kamelian	SOA	20	-	11	< 2	7	30 x 12.7 x 8.4
Kamelian	SOA	15	-	12	< 2	7	30 x 12.7 x 8.4

Table 6-10: Optical Amplifier Characteristics.

6.6 Photodiodes (PDs)

Photodiodes (PDs) are available from a wide variety of vendors and are designed for a variety of applications. As is the case for lasers, PDs are available as bare die, packaged in TO cans with and without transimpedance amplifiers (TIAs), and in receiver and transceiver modules. Table 6-11 lists the characteristics of several vendors' PDs which are packaged in TO cans either with or without TIAs.

Vendor	Type	Respons-ivity (A/W)	Dark Current (nA)	Bias Voltage (V)	BW / Data Rate (GHz/Gbps)	Wave-length (nm)	Power (mW)
Finisar	GaAs PIN	0.5	-	3.6	9/10.7	850	115
Finisar	GaAs PIN	0.5	-	3.6	4/4.25	1310	99
JDSU	GaAs PIN	0.6	0.1	2	5/4.25	850	-
Lasermate Group	GaAs PIN	0.55	0.2	5	1.9/2.5	850	-
Lasermate Group	GaAs PIN	0.3	1	5	1.5/2.5	850	-
Thorlabs	GaAs PIN	-	0.5	-	-	800-1800	-
Thorlabs	GaAs PIN	0.87	25	-	-	700-1800	-
Thorlabs	GaAs PIN	-	200	-	-	80-1800	-
JDSU	GaAs PIN w/ TIA	-	-	3.3	12/10	850	149
Lasermate Group	GaAs PIN 1x4 array	0.55	0.1	5	1.9/2.5	850	-
Thorlabs	Si PIN	0.34	2.5	-	-	200-1100	-
Thorlabs	Si PIN	0.59	20	-	-	350-1100	-
Thorlabs	Si PIN	0.56	0.6	-	-	400-1100	-
Kyosemi	Si PIN	0.3	0.015	-	2/-	850	-

Vendor	Type	Respons-ivity (A/W)	Dark Current (nA)	Bias Voltage (V)	BW / Data Rate (GHz/Gbps)	Wave-length (nm)	Power (mW)
Mitsubishi	InGaAs PIN	0.85	0.5	5	10/10	1000-1600	-
Perkin-Elmer	InGaAs PIN	0.90/0.95	4	5	3.5/-	1300/1550	100
Perkin-Elmer	InGaAs PIN	0.90/0.95	5	5	0.75/-	1300/1550	100
Bookham	InGaAs PIN w/ TIA	0.8	10	5	9/10	1550	235
Kyosemi	InGaAs PIN w/ TIA	0.88/0.98	-	15	-/10	1310/1550	201
Luxnet	Ge PIN	0.58/0.88/0.58	95	0.5	9/-	850/1310/1550	-
Univ. of Tokyo	Ge PIN on SOI	0.24	-	-	-	1550	-
Bookham	APD w/ TIA	-	100	-	7/10.7	1550	350
JDSU	APD	0.85	1	-	3/2.5	1550	-
NSG	InGaAs APD	0.8	50	-	1.8/2.5	1000-1625	-
Perkin-Elmer	InGaAs APD	8.4/9.4	6	40-90	2/-	1300/1550	-

Table 6-11: Photodiode Characteristics.

6.7 Connectors

The major types of connectors used in fiber optic communications were presented in Section 4.

6.8 Array Waveguide Gratings (AWGs)

Array waveguide gratings (AWGs) are available from several vendors. These devices are critical for wavelength multiplexing and de-multiplexing in WDM-based systems. Table 6-12 lists characteristics of AWGs from five different vendors. These AWGs are based on silicon PLC technology. Several of them are athermal, meaning they do not require temperature control, while the others have internal temperature control, typically thermoelectric coolers (TECs). All of them except Agilecom state that their AWGs are compliant with Telcordia GR-1209 and GR-1221 standards.

Vendor	Number of Channels	Channel Spacing (GHz)	Insertion Loss (dB)	Loss Uniformity (dB)	Adjacent Channel Crosstalk (dB)	Temp Controller	Size (mm^3)
Agilecom	≤ 48	100	6	1.5	23	Athermal	120 x 70 x 13
Agilecom	≤ 48	100	5.5	-	23	Internal	135 x 65 x 20
Gemfire	≤ 80	50	4.5	2	25	Internal	130 x 65 x 14
Gemfire	≤ 48	100	5.5	1.5	25	Athermal	120 x 70 x 10
Gemfire	≤ 48	100	4.5	1.5	26	Internal	130 x 65 x 14
Gemfire	≤ 8	200	3	1	25	Athermal	100 x 50 x 10
JDSU	8	50	6	1	25	Internal	130 x 65 x 14
JDSU	40	100	5	1	30	Internal	130 x 65 x 14
JDSU	40	100	2.5	1	26	Internal	130 x 65 x 14

Vendor	Number of Channels	Channel Spacing (GHz)	Insertion Loss (dB)	Loss Uniformity (dB)	Adjacent Channel Crosstalk (dB)	Temp Controller	Size (mm^3)
NEL	$\leq$ 48	100	5.5	1.5	23	Athermal	-
NEL	$\leq$ 128	25, 50, 100, 200	5.5	1.5	-	Internal	130 x 65 x 12
Neo-Photonics	$\leq$ 40	100	6	1.5	23	Athermal	-
Neo-Photonics	$\leq$ 40	50, 100, 200	4.25	1.2	25	Internal	-

Table 6-12: Array Waveguide Grating (AWG) Characteristics.

6.9 References

[6-1] Jalali, B., M. Paniccia, and G. Reed: Silicon Photonics. IEEE Microwave Magazine, 7(3):58-68 (June) 2006.

[6-2] Fedeli, J. M., et al.: Development of Silicon Photonics Devices Using Microelectronics Tools for the Integration on Top of a CMOS Wafer. Advances in Optical Technologies, Volume 2008, Article ID 412518.

[6-3] Malinge, J.-L.: A Practical Solution to Integration: Using Silicon Photonics. OIDA Photonic Integration Forum, October 7-8, 2008, Monterey, CA.

[6-4] Mekis, A., et al.: Monolithic Integration of Photonic and Electronic Circuits in a CMOS Process. Proceedings of SPIE, Volume 6897, February 12, 2008.

[6-5] McNab, S. J., and Y. A. Vlasov: SOI 2D Photonic Crystals for Microphotonic Integrated Circuits. Holey Fibers and Photonic Crystals/Polarization Mode Dispersion/Photonics Time/Frequency Measurement and Control, 2003 Digest of the LEOS Summer Topical Meetings, pp. WA2.1/79-80, July 14-16, 2003, Vancouver, British Columbia.

[6-6] Nagarajan, R., et al.: Large-Scale Photonic Integrated Circuits. IEEE Journal of Selected Topics in Quantum Electronics, 11(1):50-65 (January/February) 2005.

7 Emerging Optical Technology

This section summarizes emerging optical and photonic technologies that are, or may become, relevant to high performance compute (HPC) systems. Although quantum computing technology is an emerging field and includes photonic supporting technology, quantum computing is too immature for consideration in HPC systems at this point in time and will not be discussed further. Two other Sections, 11 and 12, also discuss active research in specific areas; optical computing elements and technology trends in the telecommunication industry, respectively. Similarly, there is potential for overlap of this section and with Section 6, Existing Optical Technology, since there are areas of overlap between existing and emerging technologies. Where there is known overlap in subject matter with this section, reference will be made to the more specific information found in Sections 6, 11, or 12. The discussion here is organized into two broad categories: a) components, and b) subsystem design strategies.

7.1 Components

As in most other fields of engineering, optical component technology continues to evolve on many different fronts driven by ever increasing system performance requirements. Two of the biggest drivers are, of course, the demand for ever increasing bandwidth and lower cost. While it is impractical to discuss current research on all of the major optical components, given the intended scope of this document, a representative sample of on-going research that is likely to influence future component development is presented in the following subsections.

7.1.1 Miniaturization and Monolithic Integration of Optical Components (Nanophotonics)

Economically, photonic integrated circuits (PICs) are considered to be essential for driving down costs of optical subassemblies (see Section 12). Higher levels of monolithic integration allows greater on-chip functionality, higher reliability, potential for higher data rates, and most importantly, lower component and operating costs to the system providers. Monolithic integration of optical components on silicon, GaAs, and InP substrates is continuing to get more sophisticated. For an example of current technology, optical component count on fielded dense wavelength division multiplexing (DWDM) laser transmitter PICs is in the range of 50 to 100. It is predicted that future PICs would require integration of several hundred, if not thousands, of optical components. Integration on silicon is especially attractive because of the possibility of combining electrical and optical integrated circuit technologies (active and passive) on the same chip.

Recent work in Germany describes silicon-organic hybrid (SOH) devices fabricated on SOI substrates. The objective of this research is to exploit and combine the strong light confining characteristics of silicon optical waveguides with the highly non-linear optical properties of certain organic liquids to realize highly compact optical signal processing components such as high speed Mach Zehnder (MZ) modulators and de-multiplexers for optical time domain multiplexed (OTDM) signals [7-1]. Examples include a 6 mm SOH device that demultiplexes a 120 Gbps data stream to 10 Gbps and an 80 µm long MZ modulator design capable of modulating at 100 Gbps rates with a drive voltage of only 1 V. The demultiplexer exploits the four-wave mixing (3^{rd} order intermodulation) phenomenon and the strong nonlinear optical properties of the organic material.

Another example of a complex (experimental) InP-based PIC is described in reference [7-2]. It is a packet forwarding chip (see diagram in Figure 7-1) that integrates a complete optical receiver incorporating two semiconductor optical pre-amplifiers (SOAs) with photodiode, and a transmitter incorporating a distributed Bragg reflector (DBR) laser followed by two electro-absorption modulators, a passive splitter and combiner and interconnecting ridge waveguides, all integrated on a 4.0 x 0.55 mm^2 die.

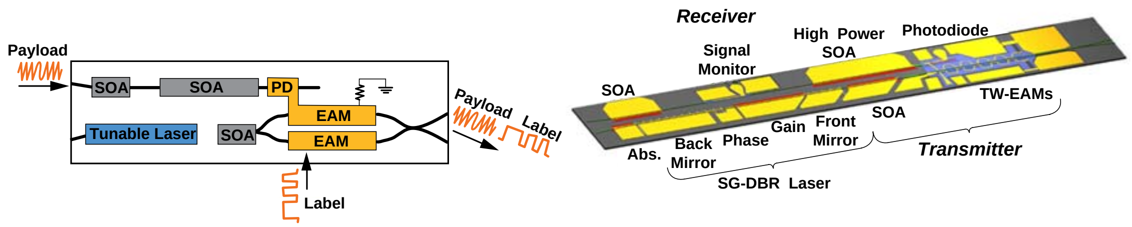

Figure 7-1:: Diagram of Field-Modulated Packet Forwarding Chip Depicting Wavelength Conversion and Label Writing Functions and Schematic of Fabricated Device Showing Integrated Components. Footprint Is 4.0 mm x 0.55 mm (From Ref. [7-2]). (40735)

Still another example of a miniature silicon-based photonic circuit is that of a non-blocking four-port bidirectional multi-wavelength message router [7-3]. The fabricated chip is approximately 0.3 mm on a side. A photomicrograph of the die is shown in Figure 7-2(a) and a color-coded example of its ability to route three different wavelengths (1538 nm, 1546 nm, and 1554 nm) incident on one port to the three other ports (or any other input/output port combination) is shown in Figure 7-2(b). It employs eight micro-ring resonators which are essentially wavelength filters that are implemented on a CMOS-compatible silicon die. Each ring is a 20 μm-diameter silica optical waveguide with cross-sectional dimensions of 450 nm x 250 nm that incorporates a heater that tunes the filter (through the thermo-optic effect) by changing the temperature and hence the refractive index (RI) of the optical waveguide. Worst case insertion loss of the signal through the router is only1.3 dB.

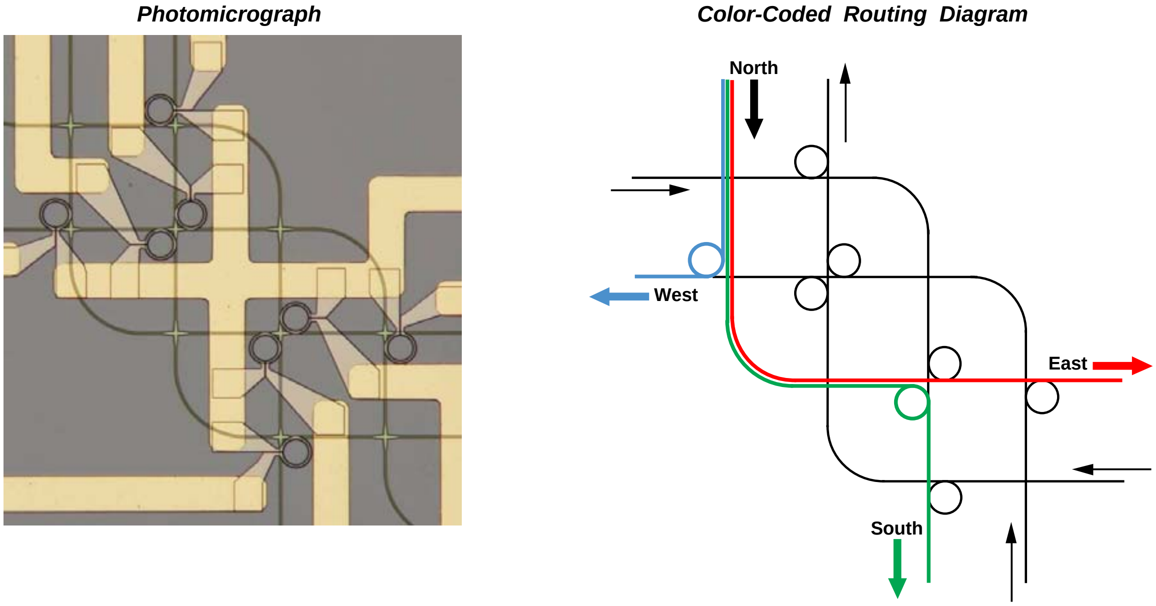

(a) **(b)**

Figure 7-2: A Non-Blocking Four-Port Bidirectional Multi-Wavelength Message Router a) Photomicrograph of Fabricated Device On A CMOS-Compatible Chip With Plan Dimensions 1/3 mm x 1/3mm and b) Color-Coded Routing Diagram Indicating Router In A Specific Configuration Routing Each Of Three Wavelengths Incident On The North Port To One Of Three Output Ports (from Ref. [7-3]). (40794)

A futuristic optical transceiver concept for intra-board high speed chip-to-chip communication such as might be used in a HPC system using CMOS-based photonic chips, is depicted in Figure 7-3 (taken from the slide presentation of reference [7-4]). The complete DWDM monolithic optical transceiver is fabricated on a CMOS chip incorporating an electro-optic (E-O) polymer optical layer that contains the optical circuit including silicon nitride waveguides, optical micro-ring resonator modulators and lateral metal-semiconductor-metal (MSM) germanium (Ge) photodetectors that are coupled to the SiN waveguides. A continuous wave (CW) multiple wavelength III-V laser array is off-chip and services all on-chip modulators. The embedded micro-ring resonators are used for wavelength filtering and high-speed wavelength-specific modulation of the light (using the E-O polymer to modulate the mode index of the SiN embedded optical waveguide). The CPU, memory, as well as the modulator drivers and amplifier electronics for the optical components are all contained in the CMOS chip directly below the optical layer.

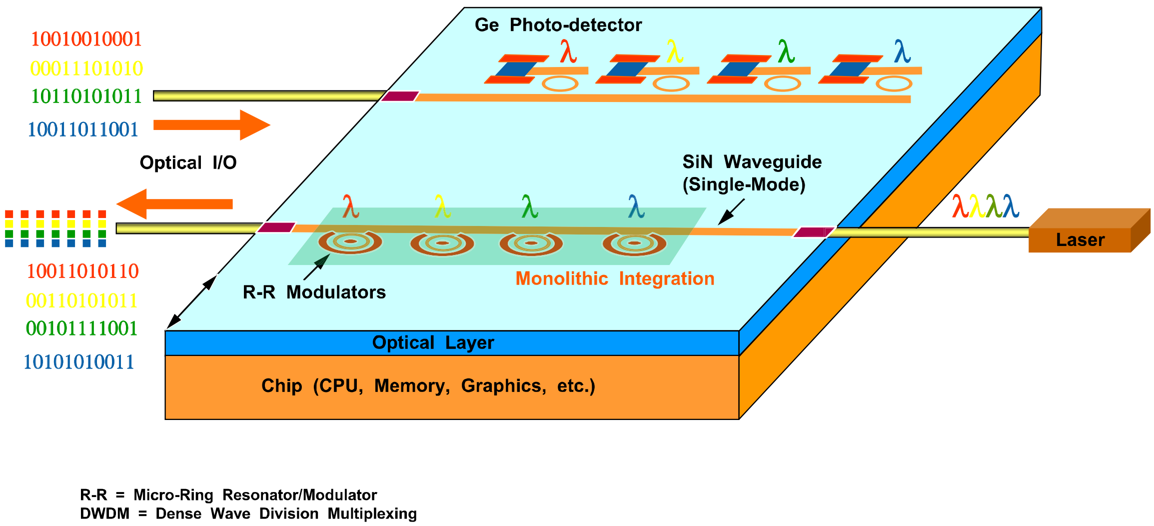

Figure 7-3: Possible Structure for CMOS-Based Monolithic Transceivers Used for Chip-To-Chip Intra-Board Optical Communication As It Might Be Applied In A Futuristic HPC system. An Electro-Optic Polymer Layer Containing The Optical Circuit Is Applied Directly On Top Of The Silicon (From Ref. [7-4]). (40795)

Finally, mention should be made of dense optical interconnects on-chip. Although there is some controversy among optical communication researchers whether intra-chip optical communications makes sense even in the distant future, there is, nevertheless, on-going research in this field. Specifically, DARPA's UNIC program seeks to explore the utility of CMOS-based intra-chip photonic interconnects for future high performance computing applications [7-5]. An argument made for the justification of this research is that the shear increase in computing power (because of continuing transistor advances) is outpacing the available bandwidth for processor-to-processor communications even at the chip level. A better balance of these two factors is needed, and on-chip optical interconnects is considered to be a potential way of achieving it. For this project, a dramatic miniaturization of optical components and power consumption of all active devices is key to the program. A power consumption target for an optical link of less than 0.01 pJ/bit has been espoused by the DARPA Program Manager.

7.1.2 Advanced Vertical Cavity Surface Emitting Lasers (VCSELs)

Vertical cavity surface emitting laser (VCSEL) technology became commercially viable towards the end of the 1990's. Since then 850 and 980-nm devices have been manufactured worldwide (with volumes at least in the 100 million range) and have been applied, primarily, to the short haul (LAN and SAN) high speed data communication markets.

Because of their many advantages over edge-emitting lasers, including lower cost, current research efforts are aimed at extending VCSEL applications to long haul systems at 1310 nm and 1550 nm including metro area networks (MANs). Furthermore, for VCSELs to be relevant to silicon planar lightwave circuits (PLCs) employing silicon ridge waveguides, for example, wavelengths greater than approximately 1000 nm are required so that the operating wavelength is beyond the silicon absorption edge and low loss optical propagation can be achieved. One of the reasons for their commercial success at 850 nm, is that GaAs-based VCSELs are relatively easy to manufacture. First, when compared to the longer wavelengths, the Bragg reflectors are more efficient (because of the high contrast in refractive index between GaAs and AlGaAs layers) and fewer layer pairs are needed for high mirror reflectivity (greater than ~99.9%) necessary for lasing. Second, the λ/4 layers are thinner because of the short wavelength and thus have lower thermal resistance than the structures for longer wavelength VCSELs. Nevertheless, the commercial "pull" for long wavelength VCSELs is driving the research for high volume manufacture of long wavelength VCSELs. Many different material approaches have been proposed and implemented in the laboratories including VCSELS on InP substrates, on GaAs substrates and hybrid implementations employing wafer-fusing of InP-based active layers with GaAs/AlGaAs mirror stacks [7-6].

Current VCSEL research is also aimed at developing devices capable of significantly higher modulation speeds. Due to reliability considerations conventional quantum well (QW) VCSELs are usually designed for modulation speeds not exceeding much beyond 10 Gbps. For higher speeds the current density can be increased, but at the expense of long term reliability due to increased device heating (modulation rate capability increases as the square root of the current density, a consequence of the basic physics of stimulated emission for QW devices). Advanced quantum dot VCSEL structures are being explored that increase the intrinsic material gain to achieve high modulation rates at reduced current densities. Other approaches include running the active layers in CW mode with low current while modulating the reflectivity of the top Bragg reflector (using, for example, the quantum-confined Stark effect to alter the index of refraction of the mirror layers). Figure 7-4 shows a diagram of this type of CW VCSEL with electro-optic distributed Bragg reflector/modulator [7-7]. Experimental devices have been demonstrated with modulation bandwidths as high as 35 GHz [7-8].

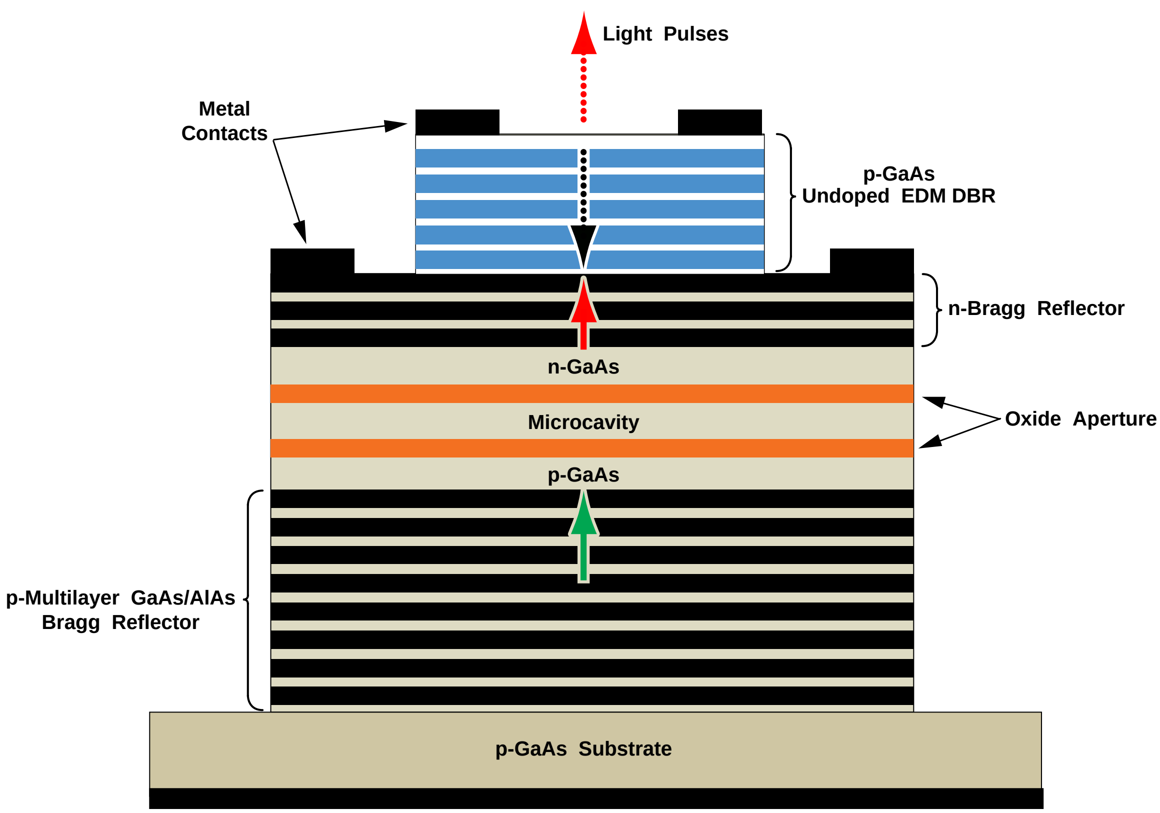

Figure 7-4: Diagram of An Advanced Continuous Wave (CW) VCSEL Structure With Electro-Optic Bragg Reflector/Modulator (From Ref. [7-7]). (40796)

7.2 Subsystem Design Strategies

This section presents general trends in optical design as applied to short interconnects and subsystems potentially applicable to HPC systems.

7.2.1 High Efficiency Silicon Photonics (Energy Harvesting)

Optical interconnects are already being used for rack-to-rack communications in HPC and data center applications. In the near future, as data rates increase, it is expected that optical links will be replacing shorter, on-board electrical interconnects as well (see next subsection). A very likely outcome in this evolution is the integration of silicon electronic integrated circuits (ICs) with optical planar lightwave circuits on silicon (e.g., monolithically integrated photonic CMOS). It is also recognized that along with the increased speeds, the advanced electronic technologies with their high transistor count and densities will require extraordinary cooling and thermal management. Together with the electronics there will be great incentive to reduce the size and power consumption of the silicon-based photonics to maintain reasonable device temperatures and long term circuit reliability. On the other hand, many of the optical functions such as modulation, wavelength conversion, switching and parametric amplification require substantial optical power (to drive the silicon material into the non-linear regime) and much of this optical power will inevitably be wasted in the form of heat. If photonics is to be an integral part of the silicon electronic/photonic platform, then it too must be designed for high efficiency and minimum heat generation.

To address this issue, on-going research is aimed not-only at increasing the photonic device efficiencies, but at actually recovering some of the wasted heat and generating electricity on-chip that in turn would power some of the electronics. At moderate to low optical intensities silicon is transparent (non-absorbing) to wavelengths above about 1000 nm. However, if the light intensity is high enough to force the onset of non-linear effects such as Raman gain and Kerr effect (required for implementing various optical functions – see Section 1 then there will also be an accompanying undesired absorption mechanism known as two-photon absorption (TPA). TPA creates heat (phonons) as well as free carriers that can absorb optical power – both effects are detrimental to the desired non-linear device operation. TPA loss can completely mask Raman gain, for example. By integrating a lateral p-n diode structure with a rib optical guide in silicon (see Figure 7-5) and properly biasing it, it is possible to sweep away the TPA-generated free carriers from the optical guide area, thus reducing waveguide loss, and at the same time produce negative dc resistance in the external circuit to generate electricity. Researchers have demonstrated energy harvesting with this two-photon photovoltaic effect (TPPV) in a 2 dB-gain Raman amplifier in silicon with a concurrent power generation of 2.5 mW [7-9]. Futuristic application of this technique in an integrated electronics/photonic IC is conceptualized in Figure 7-6.

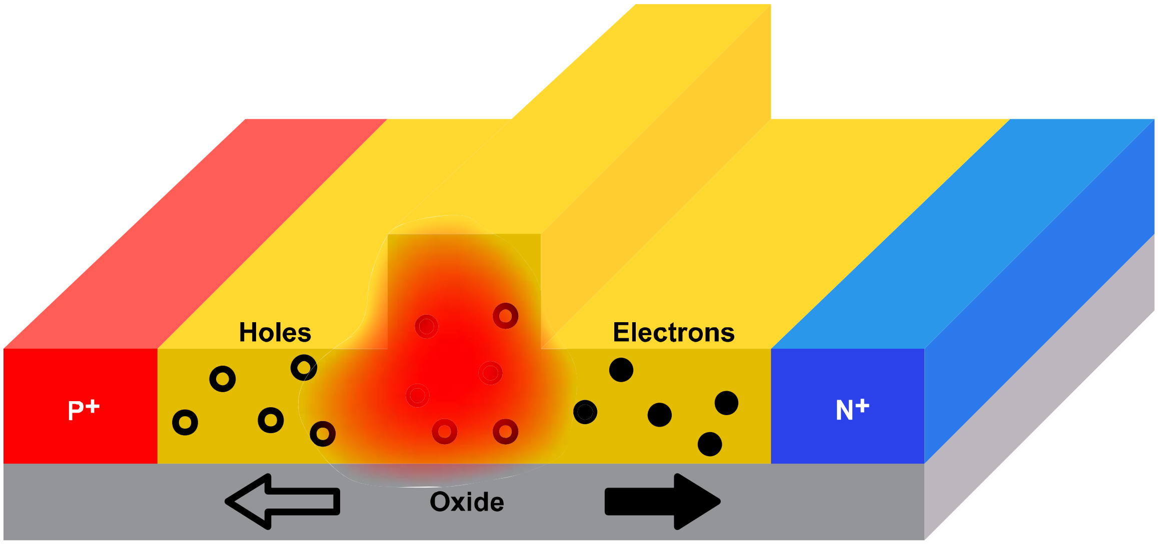

Figure 7-5: A Silicon Rib Waveguide With Integrated P-N Junction Diode That Reduces Optical Loss In The Waveguide While At The Same Time Generating Electricity Through The TPPV Effect (From Ref. [7-10]). (40741)

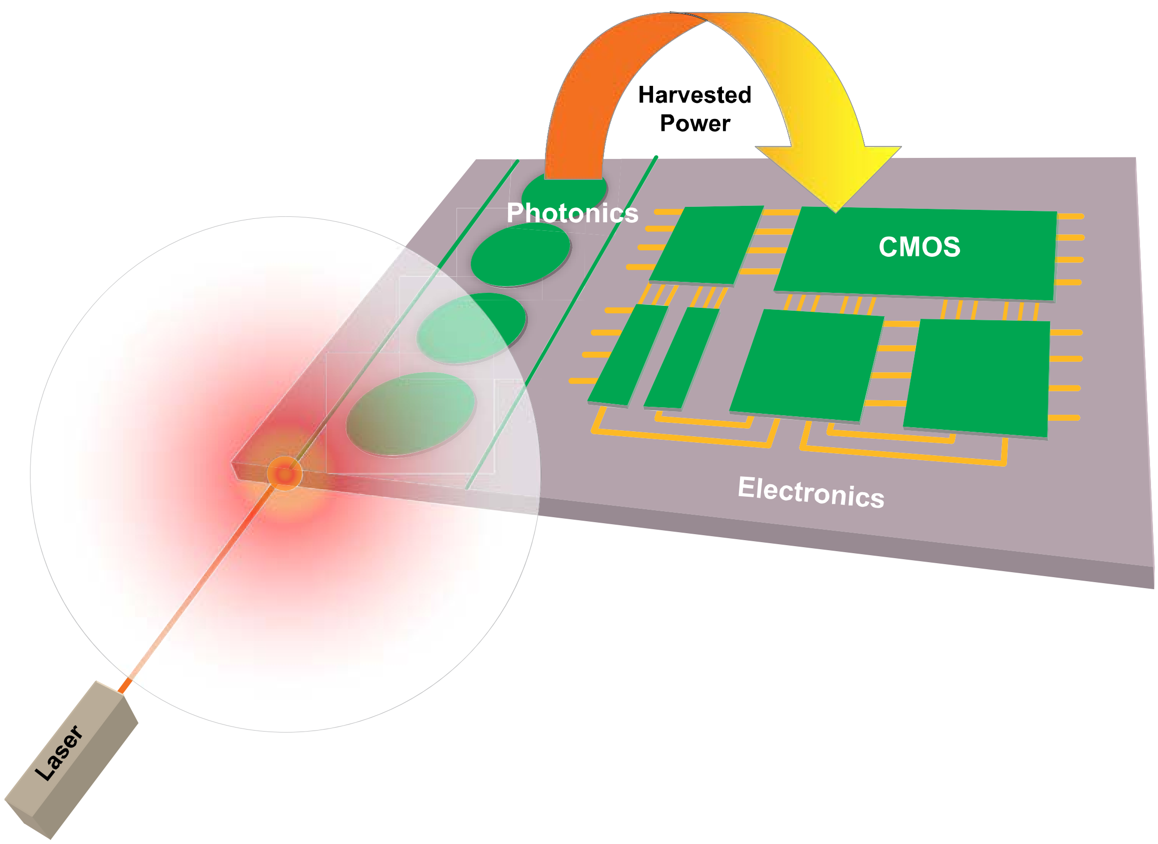

Figure 7-6: Conceptual Representation of Photonic Energy Harvesting On An Integrated Silicon Photonic/Electronic Platform (From Ref. [7-10]). (40740)

7.2.2 Optical Routing to the Electronic IC

There is considerable interest in routing optical signals not only between printed circuit boards (PCBs), but onto the board itself right up to the ASIC device. Elimination of long electrical PCB traces, together with the associated error correction and equalization circuitry, elimination of bulky O/E transceiver packages at the edge of the board, potentially smaller PCB area, and lower power consumption and cost are the usual arguments given for this general approach. In the ultimate implementation optical guides are embedded right into the PCB (see IBM's Terabus project [7-11], for example). An interim solution might be essentially a form of miniature active cable incorporated onto the board. Figure 7-7 shows a concept from Reflex Photonics [7-12]. At the edge of the board multiple channel optical connectors are used to carry the optical signals onto the board. Densely packed optical guides embedded in a flex circuit route the light to a small form factor connector and a miniature, high efficiency O/E optical sub-assembly (OSA) at the destination ASIC (an FPGA is used for the example in the drawing). The "parallel optics" flex circuit can either be in ribbon form or can have a fan-out to several different ASICS on the board. A similar approach is being pursued by Intel researchers [7-4].

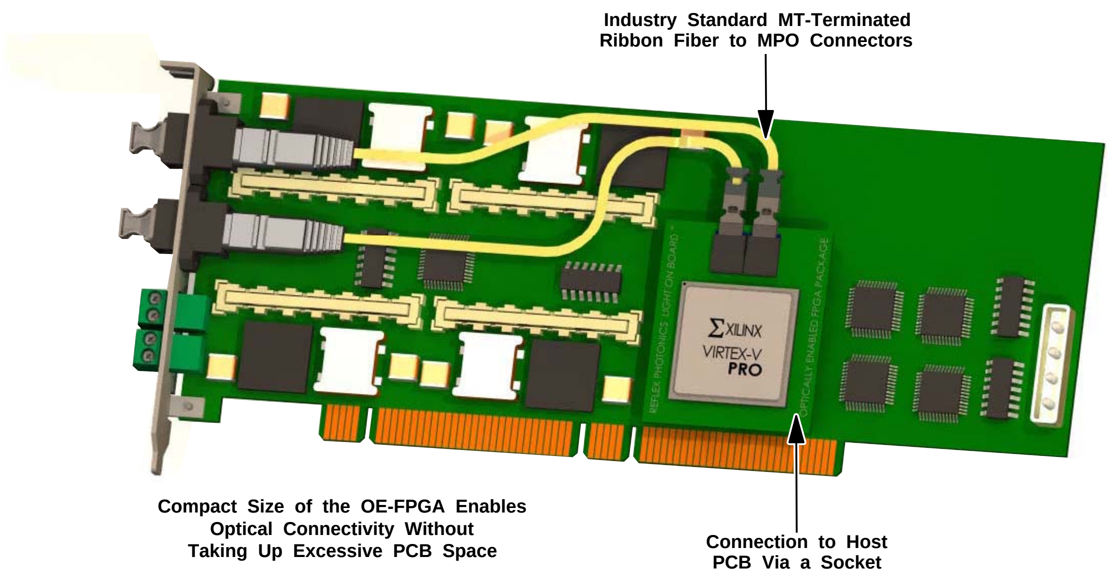

Figure 7-7: Conceptual Representation of Parallel Optics Being Routed On-Board Right Up To The Silicon ASIC (from Ref. [7-12]). (40742)

7.2.3 Optical Hybrid Integration

In contrast to monolithic integration of electrical and optical functions, there is a parallel argument among researchers in the field that considers hybrid approaches for integration. The basic idea is to employ the best material for the particular function that the device is intended for. An example in the transceiver area would be Luxtera's active cable product, "Blazar" [7-13]. Each end of the active optical cable employs a separate distributed feedback (DFB) edge-emitting laser chip (fabricated in an InP-based semiconductor) operating in continuous wave (CW) mode at 1490 nm that is mounted directly onto a CMOS silicon die. The silicon die has optical waveguides that distribute the light from the laser to four modulators. NRZ-OOK modulators and driver circuitry are implemented in CMOS. The receiver uses CMOS transimpedance amplifier circuitry that amplifies and conditions the signal from the photodetectors. The photodetectors are also InP–based and are mounted onto the silicon. The hybrid sub-assembly is a complete transceiver package incorporating four bi-directional, 10 Gbps channels that interface with 8 single-mode fibers. The "CMOS photonics" chip exploits the lower cost and maturity of CMOS technology, while the laser light is generated in continuous wave mode with a reliable InP-based DFB laser.

Another example of an optical hybrid assembly that achieves state-of-the-art performance is described in Reference [7-14]. One of the devices described is a 100 Gbps differential quadrature phase shift keyed (DQPSK) modulator. The hybrid assembly (see Figure 7-8) exploits the large electro-optic coefficient characteristics of $LiNbO_3$ and uses this substrate to fabricate precision, low loss arrayed Mach-Zehnder (MZ) phase shifters. The two adjacent silicon planar lightwave circuit (PLC) substrates contain the interconnecting optical waveguides including a thermo-optic $\pi/2$ phase shifter. One of the challenges in this subassembly is to align the optical waveguides at the two substrate interfaces. Using a combination of lateral waveguide tapering on the $LiNbO_3$ waveguides together with special angle polishing and anti-reflection coatings, and an active alignment scheme, a return loss of better than -45 dB and an insertion loss of less than 0.2 dB was achieved at the interface of the $LiNbO_3$ and silicon waveguides. The total loss for the assembly when all the MZ phase shifters are on is 7.9 dB. The assembly is packaged in a stainless steel package with overall dimensions of 9.7 cm x 1.8 cm x 0.8 cm.

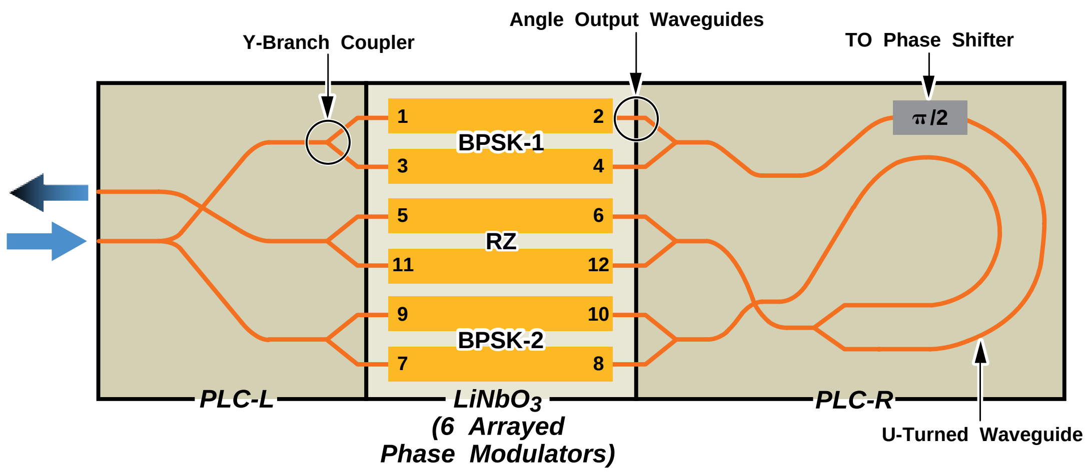

From Kaneko, A., et. al., "Compact Integrated 100 Gbps optical Modulators Using Hybrid Assembly Technique With Silica-Based PLCs and LiNb03 Devices," 2009 OSA/OFC/NFOEC, Paper No. OThN3, March 2009

Figure 7-8: 100 Gbps Optical DQPSK Modulator Assembly Fabricated On Three Separate Substrates; Two PLCs (Planar Lightwave Circuit) On Silicon Containing Optical Waveguide Couplers And A Thermo-Optic $\pi/2$ Phase Shifter And A Lithium Niobate Substrate Incorporating Low Loss Precision Mach-Zehnder Phase Shifters. An Example Of An Optical Hybrid Assembly Using Devices Fabricated On The Best Material For The Function They Are To Perform (From Ref. [7-14]). (40753)

7.3 References

[7-1] Freude, W., J. M. Brosi, P. Koos, P. Vorreau, P. Dumon, R. Baets, B. Esembeson, I. Biaggio, T. Michinobu, F. Diedrich, and J. Leuthold: Silicon-Organic Hybrid (SOH) Devices for Nonlinear Optical Signal Processing. Proceedings International Conference on Transparent Optical Networks, pp. 84-87, ICTON 2008.

[7-2] Dummer, M. M., J. Klamkin, A. Tauke-Pedretti, K. N. Nguyen, and L. A. Coldren: Field-Modulated Packet Forwarding Chips for Label-Switched Optical Routing, 2009 OSA/OFC/NFOEC, paper No. OThG3, March 2009.

[7-3] Lee, B., A. Biberman, K. Bergman, N. Sherwood-Droz, and M. Lipson: Multi-Wavelength Message Routing in a Non-Blocking Four-Port Bidirectional Switch Fabric for Silicon Photonic. 2009 OSA/OFC/NFOEC, paper No. OMJ4, March 2009.

[7-4] Young, I., E. Mohammed, J. Liao, A. Kern, S. Palermo, B. Block, M. Reshotko, and P. Chang: Optical I/O Technology for Tera-Scale Computing. IEEE International Solid-State Circuits Conference - Digest of Technical Papers, pp. 468-469, (and corresponding slide presentation), February 8-12, 2009.

[7-5] DARPA MTO Website: http://www.darpa.mil/MTO/Programs/unic/index.html

[7-6] Li, H., and K. Iga (Editors): Vertical-Cavity Surface-Emitting Laser Devices, Chapter 12. Springer-Verlag, Berlin Heidelberg, New York, 2002.

[7-7] Ledentsov, N.: Dots Speed Datacom VCSELs. Compound Semiconductor, 15(4):25-27 (June) 2009.

[7-8] Ledentsov, N.: Novel Concepts for High Speed VCSELs. 2009 OSA/OFC/NFOEC, paper No. OTuK6, March 2009.

[7-9] Fathpour, S., K. Tsia, and B. Jalali: Two-Photon Photovoltaic Effect in Silicon. IEEE Journal of Quantum Electronics, 43:1211–1217 (December) 2007.

[7-10] Jalali, B., S. Fathpour, and K. Tsia: Green Silicon Photonics. Optics and Photonics News, OPN, 20(6):18-23 (June) 2009.

[7-11] Doany, F. E., et al.: Terabus: A 160-Gb/s Bidirectional Board-Level Optical Data Bus. Proceedings of the 20th Annual Meeting of the IEEE Lasers and Electro-Optics Society, pp. 545-546, October 2007.

[7-12] Reflex Photonics Presentation to Mayo, December 12, 2008.

[7-13] Luxtera Website: http://www.luxtera.com/multirate-4x10g-optical-active-cable.html

[7-14] Kaneko, A., H. Yamazaki, and T. Yamada: Compact Integrated 100 Gb/s Optical Modulators Using Hybrid Assembly Technique With Silica-Based PLCs and $LiNbO_3$ Devices. 2009 OSA/OFC/NFOEC, paper No. OThN3, March 20, 2009.

8 Key Optical Issues for System Architects and Technologists

Throughout our evaluation of the application of optical links as applied to high performance compute (HPC) digital system designs we have looked for optical attributes and capabilities that may be of interest for HPC system architects. This section represents a summary of information and material collected that in many cases has been discussed elsewhere within this document, but for this section is either re-organized to highlight the optical issues and advantages of their use, or to explore the specific topic of interest to greater detail than previously discussed. Highlighted in Table 8-1 below is an overview of the information covered in this section.

Section #		**Issue**	
8.1		**Drivers For the Insertion of Digital Optical Communications into HPCs**	
	8.1.1	**Bandwidth-Distance Product Favors Optics**	√
	8.1.2	**Interconnect Cost**	√
8.2		**Optical Components with Unique Signal Distribution Characteristics**	
	8.2.1	**WDM**	√
	8.2.2	**Star Coupler**	√
	8.2.3	**One-to-All (Signal Fanout)**	√
	8.2.4	**In-Plane Optical Crossovers**	√
	8.2.5	**Summary of Optical Components with Unique Signal Distribution Characteristics**	√
8.3		**Power Efficiency**	√
8.4		**Miscellaneous Optics Issues**	
	8.4.1	**Optical Clock Distributions**	√

Section #		Issue	
	8.4.2	**Multimode (MM) vs. Single-mode (SM) Optics, Optical Sources and Wavelengths**	√
5.1		**Reliability of Optical Links**	**Covered in detail in Section 5**

Table 8-1: Summary of Key Optical Issues for System Architects and Technologists Covered in This Section.

8.1 Issues That May Drive the Insertion of Digital Optical Communications into High Performance Compute (HPC) Systems

8.1.1 Bandwidth-Distance Product, B*d

There are a number of leading indicators suggesting that high speed electrical I/O signaling may be reaching some of the performance limits of its ultimate high speed digital communication capabilities. Optical communications appears to offer the ability to potentially extend these limits, primarily due to significantly greater signaling bandwidths. For electrical signaling it appears that a "hard limit" for raw data rates may exist in the data rate range between 18 and 44 Gbps, depending upon the system components, the transmission line dielectric materials and the respective interconnecting line lengths. This limit is at the point where conventional channel compensation methodologies appear to be no longer capable of extending either the line length, the data rate, or both, especially at reasonable power consumption levels.

One key metric used for evaluating digital communications interconnect constructs is known as the Bandwidth-Distance Product (B*d). As noted earlier in Section 5, this metric may be used as a first level approximation to compare electrical and optical communication channels. An example of the electrical communication link length and data rate limitations for various cost effective interconnect constructs is summarized by Intel Corp [8-1], and illustrated below in Figure 8-1. If we compare the three link constructs shown, the B*d value for electrical interconnect links is within the range of 500 to 850 Gpbs-cm. Additional B*d limits are reported by other groups within Intel [8-2] in which this limit appears to be near 244 Gbps-cm for shorter length single-board channels. For an FR-4-based PCB construct with longer length and two backplane connectors this value reduces even further to 137 Gbps-cm. For a relative perspective, it is interesting to note that the B*d product for RG-58 Coax cable is reported to be on the order of 30 GHz-cm, or approximately 60 Gbps-cm for NRZ transmissions[8-3].

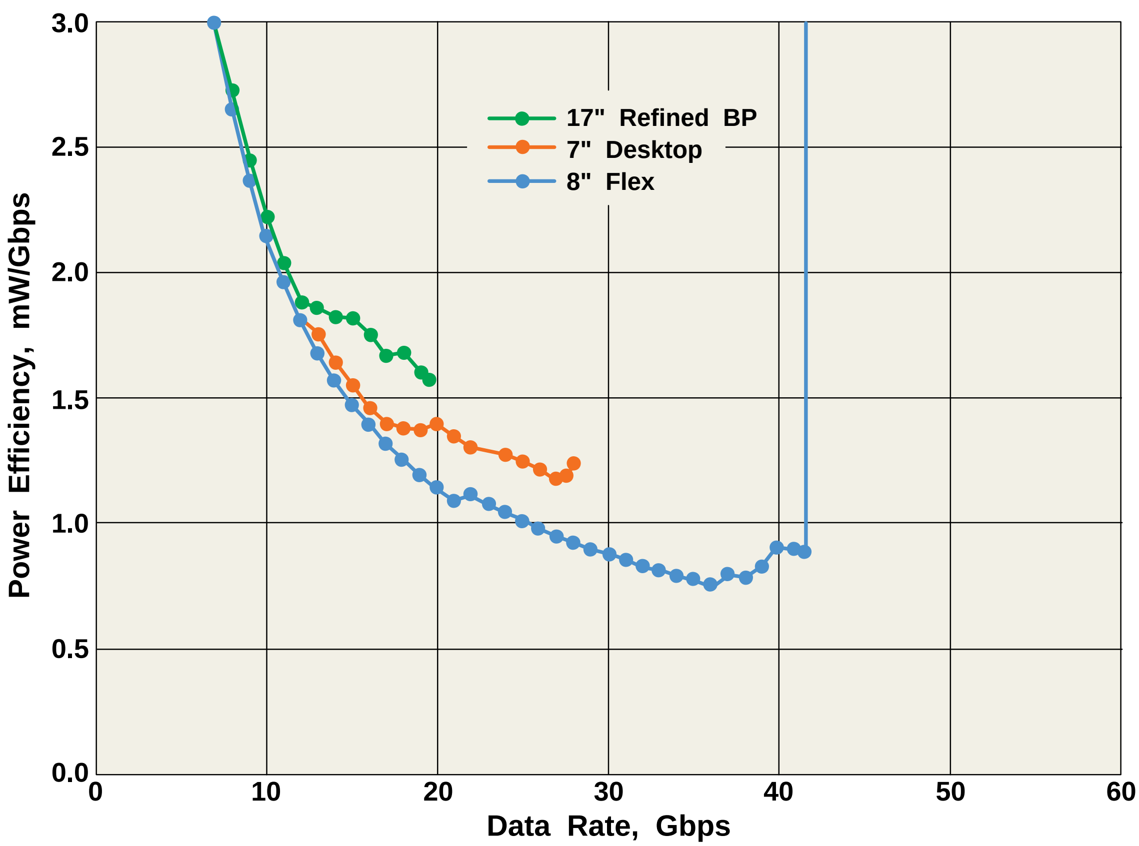

From: Young, I., et. al.: Optical I/O Technology for Tera-Scale Computing. From Slide Presentation, IEEE International Solid-State Circuits Conference February 2009.

Figure 8-1: Measured Electrical SerDes Data Rate Limitations and Power Efficiency for Three Representative Electrical Channel Constructs (For Each Data Set the Rightmost Data Point Indicates the Limiting Data Rate). (40841)

Two other references were examined for the range of values expected for electrical and optical link bandwidth-distance products. They are shown below from research investigations performed by IBM in 2008 in Figure 8-2 and Figure 8-3 [8-4]. Figure 8-2 shows the current state of the art limits for electrical signaling which when converted to B*d is approximately 810 Gbps-cm

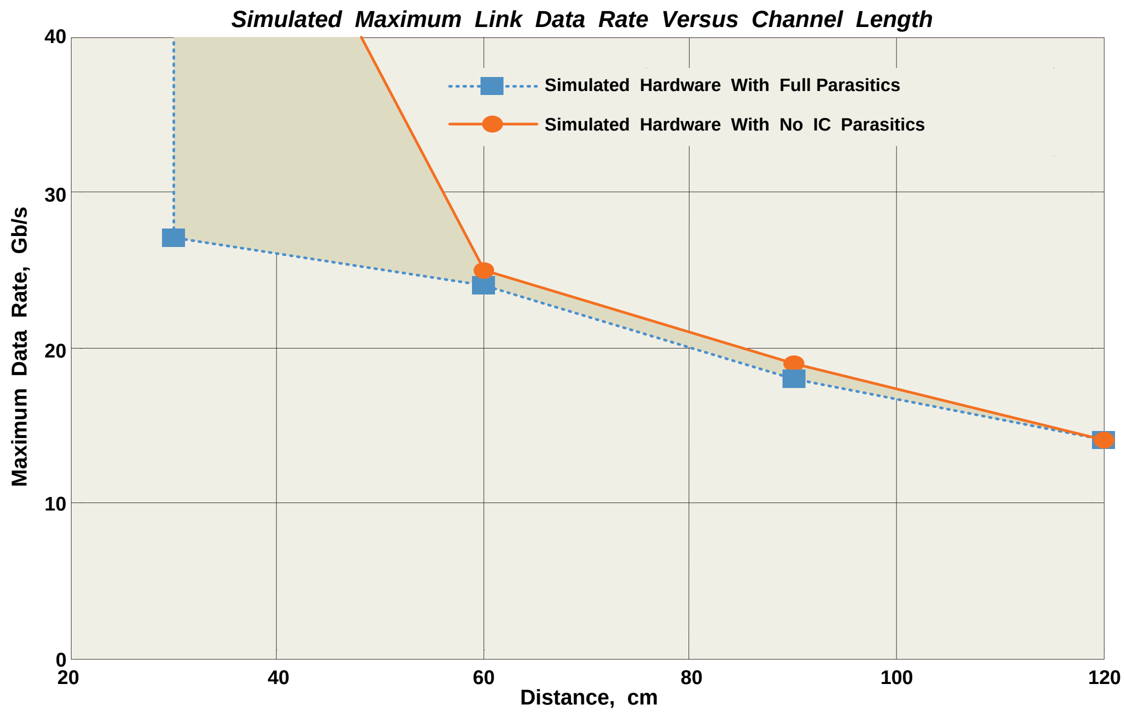

Source: Pepeljugoski, P., M. Ritter, J. Kash, F. Doany, C. Schow, Y. Kwark, L. Shaw, D. Kam, X. Gu, C. Baks: Comparison of Bandwidth Limits for On-Card Electrical and Optical Interconnects for 100 Gb/s and Beyond. Proceedings of SPIE, Volume 6897, 689701, 2008.

Figure 8-2: Maximum Simulated Electrical SerDes Data Rates. (23485)

The B*d product for optical waveguides is also expected to have a range of values depending upon the optical transmission mode as well as the dispersion characteristics of the waveguide material system in use. A range of typical optical fiber values reported for B*d is listed in Table 8-2 below [8-3].

	Single-Mode	**Multimode Graded-Index Fiber**	**Multimode Step-Index Fiber**
Source	Edge-Emitting Laser	Laser or LED	Laser or LED
Application	Under sea cables	Intercity trunks	Data Link
Splicing	Very difficult	Difficult	Relatively easy
Cost	Medium	Most expensive	Least expensive
Bandwidth-Distance product (in Gbps-cm for NRZ)	$> 6x10^5$	$6x10^5 \geq BW\ x\ Dist \geq 4x10^4$	$<4x10^4$

Table 8-2: Summary of Optical Waveguide Modes and Bandwidth-Distance Product Supported.

Figure 8-3 shows a graph of the B*d limits for multimode optical signaling across the Terabus polymer waveguide embedded within a PCB "hybrid" stackup to be in the range of 1500-3000 Gbps-cm, approaching 6000 Gbps-cm for ideal channel limits.

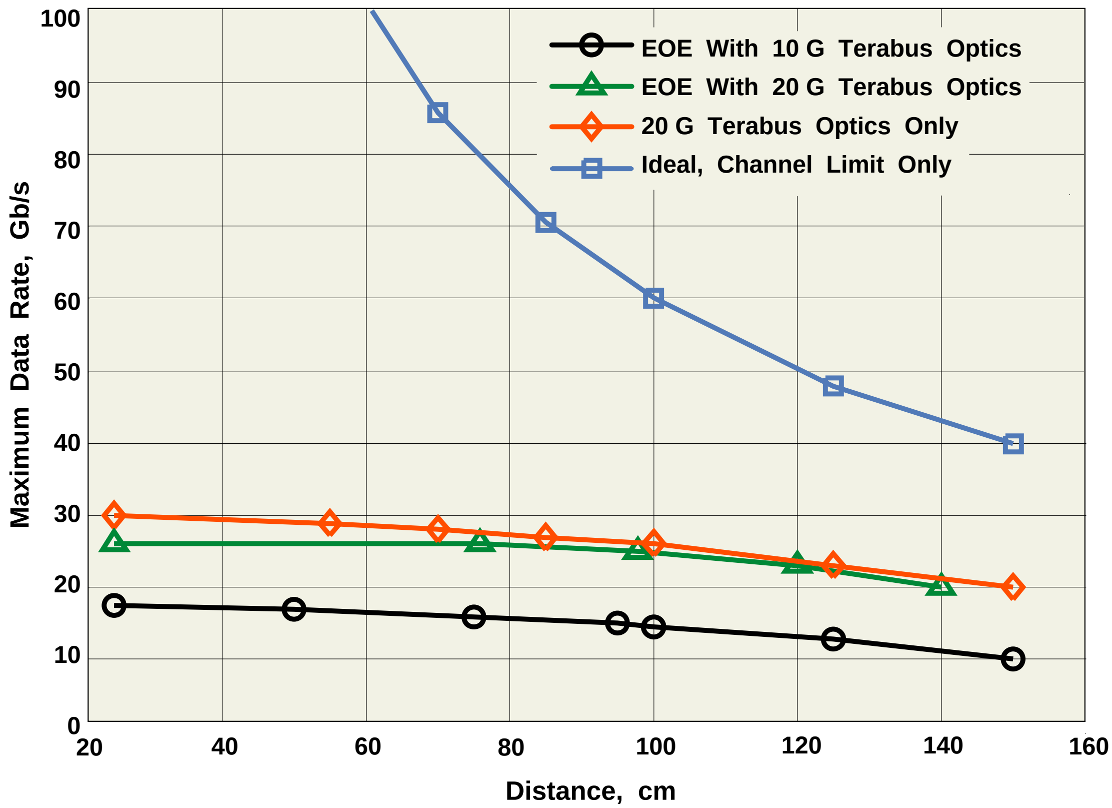

Source: Pepeljugoski, P., M. Ritter, J. Kash, F. Doany, C. Schow, Y. Kwark, L. Shaw, D. Kam, X. Gu, C. Baks: Comparison of Bandwidth Limits for On-Card Electrical and Optical Interconnects for 100 Gb/s and Beyond. Proceedings of SPIE, Volume 6897, 689701, 2008.

Figure 8-3: Maximum Multimode Polymer Optical Waveguide Bandwidth-Distance Product As Indicated By Rightmost Data Point Demonstrated By IBM Research Laboratories. (23631)

To visualize how the B*d product may be expected to affect an idealistic 36-node "all-to-all" HPC system design, we assumed a nominal value of 500 Gbps-cm for an all-electrical link. Consequently, any links exceeding 500 Gbps-cm (e.g., the longest reaches in the system) may indicate the transition to an optical solution must be made. In this manner we can use the B*d criterion to look at the percentage of optical links required within the system. The result is shown in Figure 8-4 based on processor boards with a 6X6 array of processing elements, using a manhattan routing distance of approximately 10.2 cm (4") for component-to-component pitch. The summary of the resultant link lengths exhibits a distribution that shows the percentage of interconnecting nets that may be required to be optical links with data rates of 5, 10 and 20 Gbps. For the slower rate of 5 Gbps, nearly all of the communication links can be satisfied with an all-electrical solution (only 4 links out of a total of 1260 exceeded the 500 Gbps-cm criterion), and for all intents and purposes the percentage of optical links is zero. On the other hand, at 20

Gbps, 83% of the links exceed the 500 Gbps-cm limit and required an optical link implementation.

- **Assumptions in calculation of approximate length of electrical links:**
  - **Uniform array of nodes with 4" (10.2 cm) pitch in x- and y-directions**
  - **Manhattan routing (lines are parallel to either x-axis or y-axis)**
  - **Does not include links to switches**
  - **Processors do not talk to themselves (no loopback)**
  - **B/W-distance product is based on 500 Gbps-cm estimate**

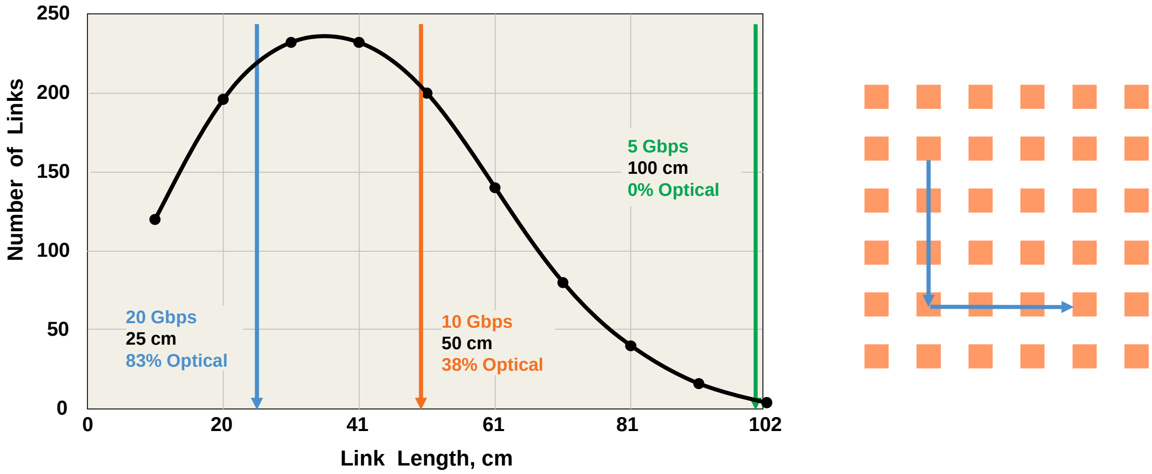

Figure 8-4: Summary Showing The Distribution of Node-To-Node Link Lengths For An Idealistic 6X6 Array Of Processing Elements and the Percentage of Links Requiring an Optical Implementation for Three Different Data Rates (Assuming An Electrical Link B*d imitation of 500 Gbps-cm). (40937)

As silicon IC technology continues to evolve along the Moore's Law trajectory, interconnect signaling bandwidth is quickly becoming a system architectural "bottleneck". As discussed earlier in section 5, we also anticipate the transitions onto high speed optical communication links within HPCs to be gradual. Table 8-3 shows a projection of where these transitions from electrical, to multimode optical, and finally to single-mode optical may occur roughly based upon the assumption that "interconnect performance demand" for B*d will double every 2 years.

Year	Electrical Signaling	B*d (Gbps-cm)	Multimode Optical Hybrid PCB Polymer Waveguide Signaling	B*d (Gbps-cm)	Multimode Optical Flex Waveguide Signaling	B*d (Gbps-cm)	Singlemode Optical Embedded Fiber Waveguide Signaling	B*d (Gbps-cm)
2009	6 Gbps	500						
2011	12 Gbps	1000						
2013			24 Gbps	2000				
2015			48 Gbps	4000				
2017					96 Gbps	8000		
2023					768 Gbps	64000		
2029					4544 Gbps	512000		
2031 and Beyond							9088+ Gbps	1024000+ Gbps

TIME

BANDWIDTH-DISTANCE PERFORMANCE

Table 8-3: Summary Of Estimated I/O Signaling Data Rate Requirements And The Respective B/W Distance Products Required To Support. (40981)

8.1.2 Optical Interconnect Cost

Many of the current optical (rack-to-rack) applications present in today's HPC designs are dominated by multimode optical links driven with VCSEL sources, and we believe the prime motive behind this current trend is cost. Current cost for DWDM links implemented in single-mode optics are estimated to be in the range of 10-50 $/Gb/s, while the cost of multimode links with VCSEL sources is significantly lower. Quotations provided to Mayo for low volume VCSEL transceivers have been in the range of 8-16 $/Gb/s for a 10 Gb/s transceiver and 25 $/Gb/s for a 5 Gb/s transceiver configured as a multi-channel active optical cable. These figures are expected to reduce again for high volume manufacturers to 1-5 $/Gb/s [8-5]. The cost advantage may widen even further with projected costs of 0.02 $/Gb/s for multimode optical links implemented in future designs containing multiple dense parallel optical channels in very high volumes [8-6].

Discussions with research and design teams from both IBM and Hewlett Packard [8-7] have indicated that the cost metric associated with the insertion of any new technology is the single most important metric in determining whether optics will appear for communication links within designs for personal computers, work-stations and servers. A cost-driven “cross-over” zone for electrical to optical links currently appears to be near the range of 1 $/Gb/s.

8.2 Signal Distribution Characteristics Unique to Optics Applications

While evaluating optical applications for HPCs we looked for optical components that demonstrated operating capabilities beyond the solutions available within the electrical signaling environment. Some of the elements and capabilities that stood out as unique are summarized below.

8.2.1 Wavelength Division Multiplexing (WDM)

Wavelength division multiplexing (WDM) allows the designer to expand the effective link bandwidth by multiplexing multiple optical signals with unique center wavelengths (also referred to as "colors") onto the same optical waveguide. There are two common forms of WDM, namely Coarse WDM (CWDM) where the wavelengths are spaced relatively far apart, and Dense WDM (DWDM) where the wavelengths are spaced more closely together. Current state-of-the-art (SOA) CWDM channel spacing is 20 nm with near-term plans to migrate product offerings to 10 nm channel spacing. SOA channel spacing for DWDM components are presently at 50 GHz, with near-term plans to migrate to 25 GHz and long-term plans heading toward 12.5 GHz channel spacing.

The major difference is the tradeoff between the significantly higher power per link required to support the thermoelectric coolers (TEC) required to stabilize the tightly controlled wavelength spacing associated with DWDM versus the elimination of this cooling overhead (potentially >50% of the entire link power budget) in exchange for fewer parallel channels with CWDM. In principal, nearly all of the DWDM based optical links we have encountered during the course of this study were based on single-mode optics with the most common center wavelength of 1550 nm typically found in long-haul telecommunication systems. Most of the CWDM based designs were also implemented at 1550 nm wavelengths, although it is not difficult to visualize CWDM designs centered at 1310 nm using SM optical waveguides and/or designs implemented at 850 nm with VCSEL sources and MM waveguides. For future HPC designs that incorporate silicon PLCs (planar lightwave circuits) with single-mode waveguides we may see architectures evolve where remote CW optical sources are used and the light is distributed throughout the system to individual modulators at the processing nodes. This approach has potentially many advantages including improved access to and maintainability of the laser sources.

8.2.2 All-to-All (Non Blocking Crossbar) STAR COUPLER

Most high speed serial electrical communications are limited exclusively to a point-to-point communication network with single transmitter and receiver elements. One of the unique elements available within an optical signal distribution is known as a star coupler, where multiple optical inputs (potentially of differing wavelengths) can be combined through the input ports and then the combined signal is available at all of the output ports. Technically, this optical device does have an electrical communications equivalent made of multiple links of coaxial cable, also known as a star coupler, developed by Digital Equipment Corporation (DEC) for a 70 Mbps link [8-8], although no higher performance electrical versions of the star coupler have re-emerged in the past 40 years.

Research reports include references to optical star couplers with up to 64 incoming and 64 outgoing (64X64) branches. Larger implementations will be potentially limited by the signal sub-division loss and the additional passive losses associated with the device. For example, an ideal 64 X 64 star coupler would exhibit a baseline loss of -18 dB for the splitting down of the input signal amplitudes, and would also be expected to exhibit an additional 3-5 dB of passive loss. In Mayo SPPDG labs we have characterized a 32 X 32 star coupler with a baseline signal splitting loss of -15 dB, and had an average of -4dB of additional passive signal loss. The excess signal loss is important for system designs as in many cases we are working with optical receivers with minimum input sensitivity thresholds of -20 to -25 dBm. We have demonstrated optical link signal broadcasts through this star coupler up to 10 Gbps. Most of the optical star couplers currently available commercially are designed for use within a single-mode optical waveguide distribution network.

For HPC applications, the uniformity between and the excess loss at the star coupler's outputs is important for higher processor node count applications as these applications will typically exercise the limits of optical receiver input sensitivities, or will require optical amplification for the output branches. Characterization data results for a 32 X 32 star coupler characterized by Mayo SPPDG are shown below in Figure 8-5.

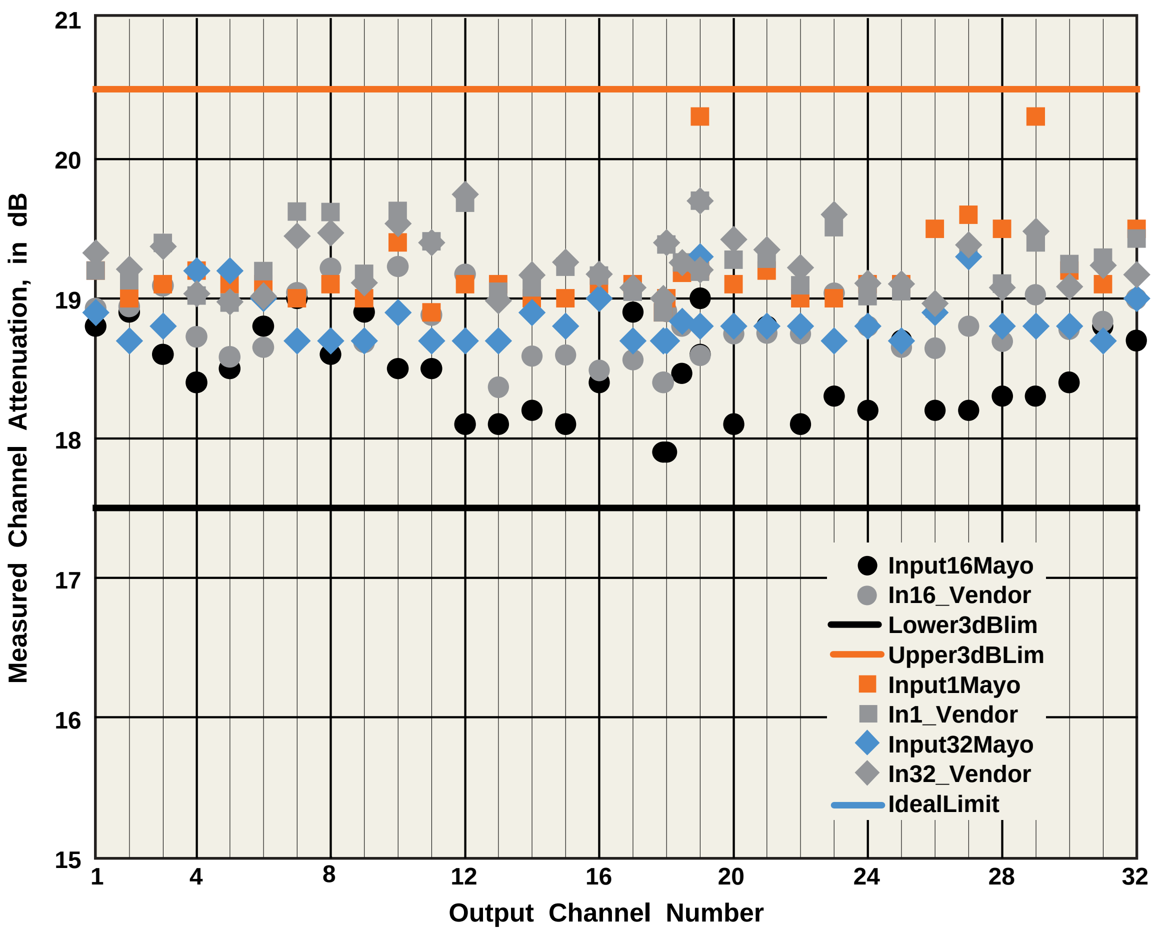

Figure 8-5: Measured Insertion Loss For A 32 X 32 Single Mode Star Coupler (Measurements Conducted At Mayo Are Compared With Those Reported By The Star Coupler Vendor). (40938)

8.2.3 One-to-All Optical Fanout/Splitter

Similar to the star coupler in principal (and potentially with respect to enabling numerically higher load fanouts) is the Optical Fanout/Splitter. In this device, a single high speed signal is broadcast out to multiple loads. This optical topology does have an electrical equivalent but as was similar to the operating bandwidth characteristics seen for the “all-to-all” topology described in section 8.2.2 the effective maximum signaling bandwidth for the analogous one-to-all electrical Fanout/Splitter is on the order of 500 MHz. The upper BW limit of the optical version of this construct has not been clearly established although it is estimated to be well beyond several hundred GHz.

An efficient (in terms of limited excess loss) design of a one-to-all signal splitter can be implemented in a planar waveguide technology with one input path, and two equal optical branch output paths resulting in a Y-shaped design. This splitter can be cascaded to achieve a tree structure consisting of multiple two-way splitters. Although they are feasible, implementations of this type are rarely seen as individual component designs, due to the difficulty associated with efficiently connecting the device to fiber inputs and outputs. Therefore this is most commonly implemented as part of a photonic integrated circuit (PIC) [8-9].

We recently encountered an alternative method for implementing optical fanouts with the hollow metal waveguide (HMG) optical channel construct methodology currently under development by HP Labs [8-10]. The HMG approach was originally conceptualized by Marcatili, et al, with Bell Labs in 1964 [8-11]. The optical waveguide is a channel (commonly implemented 150 um wide, but extensible down to a minimum width of approximately 100 um) that is cut into the substrate with a dicing saw and the three exposed sidewalls are then metallized. The fourth wall is applied with a lid to the top of the physical structure. Essentially, the light from the optical source propagates down the HMG channel, and through the use of pellicle-based channel beam-splitters a fixed percentage of the optical power (typically 100 uW for a 1 mW source) is tapped off to each load location. A detailed description of the HMG channel concept is contained in references [8-10] and [8-12]. Prototypes of this particular optical interconnect methodology have been demonstrated with up to 1:8 load fanouts. Average propagation loss for this waveguide solution is approximately 0.05 dB/cm compared to the 0.12 dB/cm measured for the IBM OptoCard Demonstration Vehicle [8-13]. One prime application cited for this optical distribution methodology is for CPU-to-memory interfaces. An additional advantage of this technology is that it works for multimode (MM) applications and also is applicable to multiple wavelength optical transmissions. Some of the concerns associated with the HMG technology are directed toward whether this capability is consistently manufacturable within a high volume, cost-sensitive environment.

8.2.4 In-Plane Optical Crossovers

One of the other interesting attributes associated with high performance optical communication channels is that optical waveguides may be "crossed-over" each other orthogonally within the same plane with relatively low insertion loss for each waveguide and with relatively high isolation between the two crossing waveguides. This optical attribute has the potential to contribute to improvements in routing and overall packaging density (assuming routing pitch is equivalent between the two), when compared to conventional copper-based electrical signaling that typically requires dedicated X-Y layer pairs that are not allowed to cross-over within the same plane and require vias to transition to other layers for 90 degree turns.

Some examples of characterized in-plane optical cross-overs are cited for short optical channels within a silicon device [8-14], where depending upon the exact geometric structures (and silicon area) used for the channel crossing the range of insertion loss per crossover can be engineered to vary between 0.16-0.65 dB. For this construct methodology, cross-coupling with respect to the channel being crossed has been measured at -30 to -40 dB.

A second reference for optical channel crossovers is based on polymer optical waveguides within a hybrid PCB construct. This has 50 µm x 20 µm cross section MM waveguide channels fabricated at 250um pitch and driven by 850 nm VCSELs. Light from the VCSELs was applied through both multimode fiber (MMF) and single-mode fiber (SMF) inputs to the PCB waveguides. Insertion loss for the crossovers excited by the MMF methodology were measured at 0.0093 dB per crossing and for the SMF methodology were measured to be 0.0058 dB per crossing with -30dB measured cross-coupling to the crossover channel [8-15].

8.2.5 Optical "Perfect Shuffle" Cross-Connects

It is advantageous to implement certain parallel processing algorithms using specific cross-connects between processors. One example is the "perfect shuffle" (PS) point-to-point cross connect scheme used in facilitating the algorithm for the fast-Fourier transform (FFT) as well as polynomial evaluation, sorting, and matrix transposition [8-16] With all optical communication links the PS can be implemented in one plane and it offers a major advantage over its electrical counterpart since electromagnetic interference, parasitic capacitance and inductive coupling issues are eliminated in the optical case. Figure 8-6 shows the cross-connect required for the case of eight processing elements [8-16]. Such PS cross-connects are available as commercial products, one form of which is implemented in the optical appliqué technology wherein the cross-connects are optical fibers embedded in a flexible polymer sheet (for additional information on the physical implementation of this optical waveguide construct, see Section 6.1.6 of this reference guide. An example of the "perfect shuffle" interconnect is shown on the left side of Figure 6-1.

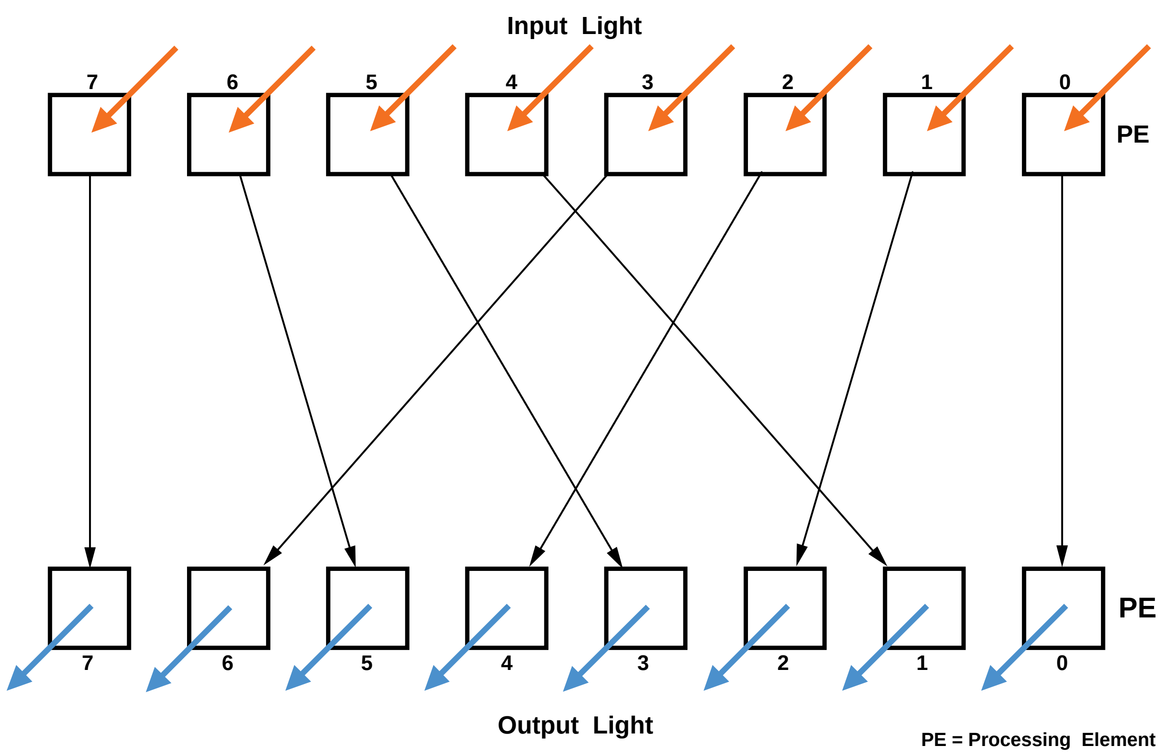

Liu, J., et al.: Implemented Optical Perfect Shuffle With a Planarized Architecture. SPIE, 3159:200-209, 1997.

Figure 8-6: Schematic Diagram of the Perfect Shuffle Cross Connect Scheme for 8 Processing Elements (PE) (Figure taken from Reference[8-20]). (40980).

8.2.6 Summary of Unique Optic Capabilities

With the exception of In-Plane Optical Crossovers, most of the listed optical components with unique signal distribution characteristics are most commonly available as single-mode COTS components. Technically it is feasible to fabricate functionally-equivalent optical components as multimode waveguide-based elements, as was demonstrated in 2003 for the λ-Connect Program [8-17], however we have not seen signs of much commercial development of any of these multimode optical components during our commercial survey and therefore would expect multimode versions of most of these components to be very expensive if available at all.

8.3 Power

The cross-over between high speed electrical signaling and high speed optical signaling in HPCs may ultimately be driven by power consumption per link. At present, it is difficult to claim that optical links are more energy efficient than their equivalent electrical counterparts, other than in the cases of long (>100 m) and medium (1-100 m) signaling lengths. A medium length example is the active optical cable (AOC) previously discussed in Subsection 5. AOCs are rapidly becoming a dominant optical insertion point for data communication applications in the 10-100 m range and beyond. One example of a successful AOC has conventional Infiniband (IB) electrical connectors on each visible end. The IB electrical signal is converted to an 850 nm multimode optical signal and is routed through an MMF to the interconnecting compute system. Power improvement for the AOC interconnect comes from not being required to repeat and re-amplify the signal for every 8 meters of cable length, as is required for the electrical analog. An IB AOC is very energy efficient with a 100 m cable containing four 10 Gbps signals requiring a total power of 2 watts, or equivalently 25 pJ/bit.

To examine other potential HPC optical insertion points based exclusively upon power, we next looked more closely at some of the high speed SerDes electrical links. Although we've reviewed reports from research teams that have proposed a fully optical SerDes link with optical arbitration for future systems [8-18], it is our belief that initial optical links will be driven by an electrical SerDes on the transmitter end and similarly received on the receiver end by conventional electrical SerDes design. To gain an understanding of the energy usage of such SerDes circuits, the energy allocation and efficiency of specific constituent elements of a typical SerDes is shown in Figure 8-7 [8-2].

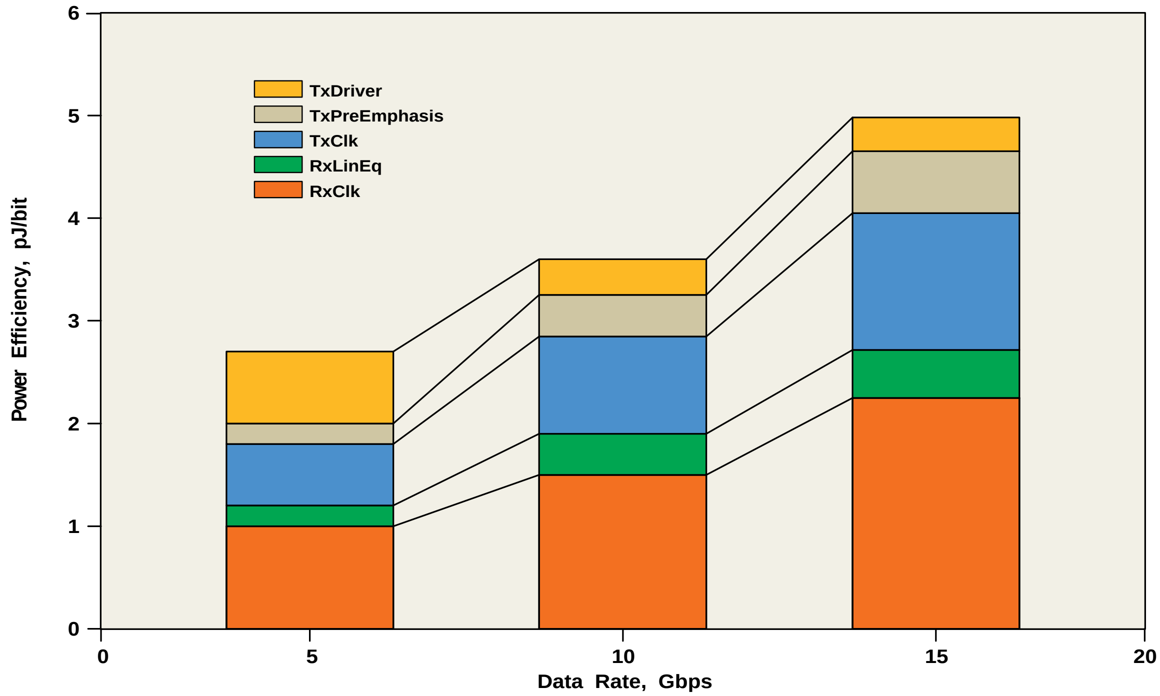

**Source: Balamurugan, G., et al.: A Scalable 5-15 Gbps, 14-75 mW Low-Power I/O Transceiver In 65 nm CMOS. IEEE Journal of Solid-State Circuits, 43(4), April, 2008.*

Figure 8-7: Electrical SerDes Link Power Efficiency Separated By Functional Block. (40839)

One power efficiency benefit associated with the use of optical links is that it is assumed that the electrical SerDes transmitter (Tx) and receiver (Rx) would be physically very close to the optical driver and receiver elements, and therefore it may be possible to gain energy efficiency through reduced Tx driver power as well as a reduction or potential elimination of channel compensation usually represented as either Tx pre-emphasis, or Rx linear equalization (LinEq). Examination of Figure 8-7 reveals that these components account for approximately 30% of the link power. Looking back to Subsection 5.1, we see that the normalized power consumption of an electrical SerDes link design is on the order of 10 pJ/bit, with the absolute "best in class" designs achieving 2.2 pJ/bit. Assuming optical-electrical replacement will occur when power parity is reached, we see optical links being required to achieve link efficiencies on the order of 1-3 pJ/bit which is 5-10X better than most of the SOA optical links in existence today.

Our conclusion is that current optical links are not competitive with electrical links based upon power and therefore, for the near-term may only be driven into HPC designs based on one of the other attributes previously discussed (B*d, cost, unique capabilities).

8.4 Miscellaneous Optical Issues and Topics for HPC System Architects

This section contains a few miscellaneous optical issues and topics suggested for consideration by system architects, but that do not logically fit elsewhere within this document.

8.4.1 Optical Clock Distributions

We conducted a search of recent journal articles and publications in the area of optical clock distributions for HPC system designs. In particular we were looking for additional information on clock phase noise and clock power dissipation. For clock phase noise we found limited content regarding the advantages inherent in optical clock distributions. It was assumed that this may have been due to the recent trends in SOA HPC system designs migrating away from single clock source designs with a limited number of clock domains and migrating toward multiple clock domains with multiple clock sources and data transmission resynchronized through the use of clock data recovery (CDR) circuitry. One literature example cited a CDR circuit design operating at 35 to 46 Gbps with rms jitter characteristics measured at less than 0.9% of the full eye pattern Unit Interval (UI). For this CDR design the measured jitter results were 226 fs RMS jitter and 1.56 ps peak-to-peak jitter measured for the recovered clock [8-19].

An additional journal article originating from Professor James Meindl's group from the Georgia Institute of Technology [8-20] utilized the 2002 International Technology Roadmap for Semiconductors (ITRS) to cite a local 19.3 GHz on-die clock distribution needed to support CPU designs associated with 32 nm minimum dimensions, with a compelling need for an on-die optical clock distribution. Unfortunately, increasing CPU clock distribution requirements stopped evolving well before the frequencies cited by ITRS were ever reached. This advancing core clock evolution stopped occurring at the introduction of the 90 nm process technology node that arrived in the 2003-2004 timeframe. This article also acknowledged the lack of "…an optical receiver capable of high-speed, low-jitter, and low-power operation…". Ensuing conversations with the optical research team at IBM Zurich also indicated that their evaluations of optical clock distributions for HPC system designs have revealed "no significant advantages or benefits" for optical clock distributions [8-5].

8.5 A Brief Comparison of Multimode Versus Single-Mode Optics for High Performance Computing (HPC) Systems

As indicated earlier in this section, we anticipate that high-speed digital data communications will evolve from the high-speed electrical signaling present today, into multimode optics on a very gradual basis. We anticipate multimode optics potentially inserted as a displacement for conventional copper interconnects at the PCB and system packaging level, first appearing as polymer waveguides within a "hybrid" PCB construct containing both multimode optical and electrical signaling channels. The second stage is envisioned to include the development of multimode waveguides with B*d limits exhibiting characteristics much closer to those associated with conventional multimode fiber waveguides, but within a dense flex waveguide construct. Ultimately we see the long-term HPC interconnect evolution heading toward a single-mode waveguide solution for digital optical communications in HPCs. Some of the characteristics driving these trends are the increased B*d and the precise signal manipulation and control available with single-mode optics. As mentioned, this evolutionary trend will likely be very gradual, in part because of the need to improve some of single-mode optic's fundamental impediments, such as cost, power, and precision waveguide alignment requirements. Below in Table 8-4 we provide an overall qualitative summary of the advantages and disadvantages of single-mode versus multimode optical characteristics when applied to HPCs.

Issue	Multimode	Single-Mode
Reliability		+[1]
Performance		+[1]
Power and Efficiency	+	
Channel Density	+ (today)	+ (future)
Distance B/W Product		+
Cost	+	
Tolerance to Mis-Alignment	+	
Efficient Couplers Today?		+

Table 8-4: Brief Comparison of Advantages and Disadvantages of Single-Mode Optics and Multimode Optics for HPC Application.

8.6 References

[8-1] Young, I., et al.: Optical I/O Technology in Tera-scale Computing. IEEE International Solid State Circuits Conference (ISSCC), pp. 468-469, 469a, February 8-12, 2009, San Francisco, CA.

[8-2] Balamurugan, G., et al.: A Scalable 5-15 Gbps, 14-75 mW Low-Power I/O Transceiver in 65 nm CMOS. IEEE Journal of Solid State Circuits, 43(4):1010-1019 (April) 2008.

[8-3] Thiagarajan Viswanathan: Telecommunication Switching Systems and Networks. New Delhi: Prentice-Hall of India, 1998, p. 234.

[8-4] Pepeljugoski, P., et al.: Comparison of Bandwidth Limits for On-Card Electrical and Optical Interconnects for 100 Gb/s and Beyond. Proceedings of SPIE, Optical Interconnect Technologies II: Joint Session With Conference, 6897:68970I(7 pages), January 21, 2008, San Jose, CA.

[8-5] Offrein, B., et al.: WDM for HPC, Consideration for the discussion with Mayo Clinic; Presentation Material from IBM Zurich Research Laboratories Presented September 16, 2009.

[8-6] Trezza, J., et al.: Parallel Optical Interconnects for Enterprise Class Server Clusters: Needs and Technology Solutions. IEEE Optical Communications, pp. s36 – s42, February 2003.

[8-7] Astfalk, G.: Chief Scientist, Hewlett Packard Laboratories (HP Labs), Private communication conducted August 16, 2009.

[8-8] Wikipedia definition of "Star Coupler": http://en.wikipedia.org/wiki/Star_coupler

[8-9] Dutton, H. J. R.: Understanding Optical Communications, IBM Redbook 1998. http://www.redbooks.ibm.com/redbooks/pdfs/sg245230.pdf

[8-10] Tan, M., et al.: A High-Speed Optical Multi-Drop Bus For Computer Interconnections. Applied Physics A - Material Science & Processing, Springer-Verlag, vol. 95, no. 4, pp. 945-953, June 2009.

[8-11] Marcatili, A., et al.: Hollow Metallic And Dielectric Waveguides For Long Distance Optical Transmission And Lasers. Bell Syst. Tech. Journal, 1964.

[8-12] Bicknell, R., et al.: Fabrication and characterization of hollow metal waveguides for optical interconnect applications. Applied Physics A - Material Science & Processing, Springer-Verlag, 95(4):1059-1066 (June) 2009.

[8-13] Doany, F. E., et al.: 160 Gb/s Bidirectional Polymer-Waveguide Board-Level Optical Interconnects Using CMO-Based Transceivers. IEEE Transactions on Advanced Packaging, 32(2):345-359 (May) 2009.

[8-14] Poon, A.W., et al.: Cascaded Microresonator-Based Matrix Switch for Silicon On-Chip Optical Interconnection. Proceedings of the IEEE, 97(7):1216-1238 (July) 2009.

[8-15] Bamiedakis, J., et al.: Low Loss and Low Crosstalk Multimode Polymer Waveguide Crossings for High-Speed Optical Interconnects. Conference on Lasers and Electro-Optics, 2007. CLEO 2007, May, 2007, pp. 1-2.

[8-16] Liu, J., D. Sun, C. Zhao, and R. Chen: Implemented Optical Perfect Shuffle With A Planarized Architecture. SPIE, 3159:200-209, 1997.

[8-17] Patel, R., et al.: Multiwavelength Parallel Optical Interconnects for Massively Parallel Processing. IEEE Journal of Selected Topic In Quantum Electronics, 9(2):657-666 (March/April) 2003.

[8-18] Beausoleil, R., et. al.: A Nanophotonic Interconnect for High-Performance Many-Core Computation. IEEE LEOS Newsletter, pp. 15-22 (June) 2008.

[8-19] Noguchi, H., et al.: A 35-to-46-Gb/s Ultra-Low Jitter Clock and Data Recovery Circuit for Optical Fiber Transmission Systems. Compound Semiconductor Integrated Circuit Symposium, CSIC 2007, October, 2007, pp. 1-4.

[8-20] Mule, A., et al.: Electrical and Optical Clock Distribution Networks for Gigascale Microprocessors. IEEE Transactions on Very Large Scale Integration (VLSI) Systems, 10(5):582-594 (October) 2002.

9 Key Questions to Ask Optical Suppliers

This section includes suggested guidelines for discussions with vendors of optical components for high speed data communication including components for high performance computing (HPC) applications. For component vendors these guidelines are summarized in Table 9-1 with both generic, and business related questions that might be asked of them. Table 9-2 is organized according to major optical component types and lists some basic technical questions related to device specifications.

A second set of guidelines is suggested for use in discussions with HPC system integrators. For this case, the set of questions and discussion points are somewhat distinct from those involving component vendors. The HPC system integrator questions are listed separately in Table 9-3. It is understood that many of the questions listed in these tables would most likely be considered sensitive by the party being queried. There would likely be a reluctance to answer them in any amount of detail.

Types of Products and Their Characteristics	Business and Financial Character	R&D Activities and Future Product Directions
What is the core product line, or primary value-added service (vertically integrated, or focus on specific product and/or service)?	What is the main source of revenue (product sales, venture capital, government contract support)?	What are some examples of advanced products?
Who are the end users (system integrators or subsystem integrators)?	Number of employees? Foreign - or domestic-owned? What percentage is off-shore?	What is the schedule for release of advanced products?
Is the product line mainly focused on multimode (MM) vertical-cavity surface emitting laser (VCSEL)-based (parallel optics) or single-mode edge-emitting lasers for dense wave division multiplexing (DWDM) systems?	Any product specifically for HPC application?	Any research activity addressing higher data (symbol) rates?

Types of Products and Their Characteristics	Business and Financial Character	R&D Activities and Future Product Directions
For active components what is the power consumption, or energy efficiency per link, and what is the overall conversion efficiency? Most high speed optical digital serial communication links require energy for both electrical and optical active components. Please define the dividing line between the two, and account for energy provisions for each. What specific high speed circuit elements are accounted for?	Do you see a market driven need for reducing power consumption for your products in near term?	Any programs addressing reduced power consumption?
For active products: What is the basic construct (i.e., InP photonic integrated circuit (PIC), hybrid optoelectronic IC (OEIC) on silicon, planar lightwave circuit (PLC), other)?	Willingness/ability to address modest volume product demand, i.e., for new/unique product capabilities; have you established any second source production capabilities? (e.g., for HPC)	Is there reliability data available that can be shared?
For active products: What is the highest single channel serial digital data rate you support?	Who are your main competitors?	Any consumer mass market opportunities for your products?

Table 9-1: Suggested Points of Discussion for Vendors of Optical Components.

Component Type	Key Parameters			
Lasers	Optical Wavelength [nm]; Tunability (range, resolution and speed);	Optical Output Power [mW] ; Slope efficiency [mW/mA] Threshold current [mA]	Operating temperature range [^{o}C] Reliability at temperature [FIT]	Conversion Efficiency [%]; Power required for temperature stabilization [mW] Direct modulation bandwidth [GHz]
Intensity Modulators **Phase Modulators**	Drive voltage [V] V_{π} [V] (voltage to effect 180^{o} phase change)	Extinction ratio [dB] OMA (Optical Modulation Amplitude) [mW]	Insertion Loss [dB]	Bandwidth and center wavelength [nm] or [GHz]
Multimode Optical Fiber **Single-mode Optical Fiber**	Glass or Plastic (POF) Non-dispersion-shifted fiber (NDSF); dispersion-shifted (DSF); non-zero dispersion-shifted (NZ-DSF); polarization maintaining (PM); many other specialized fiber	Core/Clad Diam. 50/125 [μm] 62.5/125 [μm] MM; 9 /125 [μm]	Bandwidth-distance product (GHz-m) Attenuation [dB/km]	Single fiber or multiple fiber ribbon

Component Type	Key Parameters			
Photodetectors	Optical Bandwidth [nm]	Responsivity [mA/mW]	Electrical bandwidth (GHz) Shunt capacitance [pF]	Optical power handling [mW]
Couplers	Power splitting loss and excess insertion loss [dB]	Directivity [dB]	Multimode (MM) or Single-mode (SM)	Number of input and output ports
Wavelength Selective Devices	Bandwidth [nm] or [GHz]	Insertion loss [dB]	Free spectral width [nm] or [GHz]	Optical return loss [dB]
Amplifiers	Type: erbium doped fiber amplifier (EDFA), erbium doped waveguide amplifier (EDWA), semiconductor optical amplifier (SOA) Power consump-tion [mW]	Optical Gain [dB]	Saturation power [mW]	Bandwidth [nm]
Transceivers	SM or MM and wavelength band	Power consumption [mW] Max data rate [Gbps]	Energy efficiency [pJ/bit]	Connector and package type (numerous)

Table 9-2: Important Characteristics and Metrics for Optical Components.

Overview Technical Questions	Business Interest Regarding Government-Sponsored Work	R&D Activities
What major technology breakthroughs are required for intra-board optical interconnects to be commercially viable?	Does the company foresee augmenting their R&D with government-sponsored contract work? How do you see this integrating with, and enhancing your product development plans?	In general terms what are your primary areas of research related to optical interconnects?
For intra-board interconnects what basic digital communication link construct is envisioned?	If yes to above question, – in what specific areas?	What do you see as the main barriers to optical system insertion at the intra-board level? What research activities are needed to overcome these barriers?
Do you envision optical interconnects (especially on-board and intra-IC links) having a major impact on HPC architecture? If so, in what areas and in what time frame?	What are the current government-sponsored programs your company is working on?	Is there any current research activity related to on-chip optical interconnects?
Do you see a potential application for optical interconnects being used in processor-to-memory communications? If so, what enabling/breakthrough capabilities do you feel optical interconnects can potentially provide?	Please describe what you see for a Total Area Market (TAM) for your product?	Are you involved with silicon photonics research? If so, can you describe, at least in general terms, specific activities?
What do you see as elements of optical infrastructure that is currently missing, and how would you go about putting that infrastructure in place to ensure a successful production environment?	Which of your divisions are mainly focused on products for the government that require optical data com technology?	Are there specific projects related to 3-D packaging of chips of different materials integrating optical and electronic functions?

Table 9-3: Suggested Points of Discussion for HPC System-Integrator Businesses.

10 Optical Standards

10.1 Introduction

During the course of our research in the area of optical data communication for high performance compute (HPC), we have encountered a wide array of optical specifications originating from multiple standards organizations. Standards encourage interoperability between multiple component sources, such as defining the dense wave division multiplexing (DWDM) wavelengths (as defined by the International Telecommunication Union (ITU)), or defining the environment that a component is required to function within (such as the Mil Specs). The telecom industry has worked to develop products and infrastructure that conform to these standards to allow the industry to work worldwide seamlessly and to reduce their infrastructure cost. The HPC environment, from a standards perspective, is related to the telecom environment as suggested in Table 10-1.

Environment	Telecom	HPC
Optical Path	Long >> km	Short < 300 m
Temperature Range	Wide	Narrow
Optical Power	High	Low
Fiber (or Optical Waveguide) Quality	High (Concern for Chromatic Dispersion, Polarization, <1 dB/km Loss, Etc.)	Low (Concern for <1 dB/cm Loss, High Packing Density, Ease of Alignment, Etc.)
Protocol Examples	Synchronous Optical Network (SONET), Ethernet	InfiniBand

Table 10-1: General Requirements Comparison Between Telecom and HPC Environments.

There are two primary guideline mechanisms within the optical industry:

- The international and national standards organizations, such as ITU,
- The industry collaborative standards, such as the Multi-Source Agreement (MSA).

10.1.1 International and National Standards Organizations

Standards make it possible for the internet to work worldwide as it provides compatibility between parts from different manufacturers. A key role that standards organizations provide is to continually update standards in a timely and organized manner as the technology evolves.

The ITU is a branch of the United Nations that works to coordinate the global telecommunication networks and services between governments and the private sector. Their goal is to create standards before the industry creates a 'de facto' standard which divides and hurts the growth of the market place.

Three of the US organizations that perform standards work include the Electronic Industries Alliance (EIA), Telecommunications Industry Association (TIA), and Telcordia.

10.1.2 Industry Collaborative Standards

Industry collaborative standards are developed under MSAs. An MSA is not an organization, such as the ITU, but a specific document that is drawn up between vendors that defines cooperation in agreeing to a set of specifications to promote the commercialization of a product design including its packaging. For example, the XFP (10 Gigabit Small Form Factor Pluggable) MSA specifies the module, cage hardware, electrical interface and the optical interface. Each vendor still applies their own expertise to the design. Typically, the MSA will incorporate industry standards in defining their specifications. Several examples of MSA documents can be found on the web. For example[10-1] is for the hot-pluggable 40 Gb/s optical transceiver and [10-2] is for a 10 Gigabit small form factor pluggable module**.**

10.1.3 Emphasis of Additional Sections

We have organized the following subsections to briefly summarize some of the relevant optical standards for HPC applications. These are: optical fibers, optical connectors, eye safety, wavelength division multiplexing (WDM) and optical packaging.

10.2 Optical Fibers

There are a large number of standards related to optical fiber associated with the US standards organizations (NEC (National Electric Code), EIA and TIA) and the international standards organizations (ITU and IEC). A full list of optical fiber standards is outside the scope of this document, so we will briefly describe a few of the major ones in this subsection to give the reader an idea of the breadth and substance of this topic. The physical attributes of optical fibers are covered in Section 2 Fiber Types and Other Transmission Media and Section 6 Existing Optical Technology.

Included here are samples of standards applicable in the US as defined by the NEC, EIA and TIA. The only mandatory standard in the US is the NEC 770 which specifies fire prevention standards for fiber-optic cables.

The EIA-TIA has a series of standards related to optical fiber which are commonly known as Fiber Optic Test Procedures (FOTP), but are numerically designated as TIA-455-x. For example, TIA-455-229 is the FOTP-229 standard. Other similar examples are given in Table 10-2 [10-3].

TIA-455-XX	FOTP-XX	Description
TIA-455-229	FOTP-229	Optical Power Handling and Damage Threshold Characterization
TIA-455-14-A	FOTP-14	Fiber Optic Shock Tests (Specified Pulse)
TIA-455-21-A	FOTP-21	Mating Durability of Fiber Optic Interconnecting Devices

Table 10-2: Abbreviated List of FOTP Standards.

There is also a set of Optical Fiber Standard Test Procedures (OFSTP), several examples of which are given in Table 10-3 [10-3].

TIA-526-XX	OFSTP-XX	Description
TIA-526- 11	OFSTP-11	Measurement of Single-Reflection Power Penalty for Fiber Optic Terminal Equipment
TIA-526-4- A	OFSTP-4	Optical Eye Pattern Measurement Procedure
TIA-526- 7	OFSTP-7	Measurement of Optical Power Loss of Installed Single-Mode Fiber Cable Plant

Table 10-3: Abbreviated List of OFSTP Standards.

Reference [10-3]**,** The Fiber Optic Association Tech Topics, provides a more complete list of FOTP and OFSTP standards. It also lists additional TIA standards, some examples of which are given in Table 10-4.

Standard	Description
TIA-458-B	Standard Optical Fiber Material Classes and Preferred Sizes
TIA-598-C	Optical Fiber Cable Color Coding
TIA-492AAAA	Detail Specification for 62.5 micron Core Diameter/125 micron Cladding Diameter Class I A Multi-node, Graded Index Optical Waveguide Fibers

Table 10-4: Abbreviated List of Additional EIA Standards.

Included here are samples of fiber optic international standards as defined by the ITU and IEC.

The ITU-T G.65X describes the optical fiber core material; examples are given in Table 10-5 [10-4], [10-5].

Standard	Description
ITU-T G.653	Dispersion-Shifted Single-Mode Optical Fiber and Cable
ITU-T G.654	Cut-Off Shifted Single-Mode Optical Fiber and Cable
ITU-T G.655	Non-Zero Dispersion-Shifted Single-Mode Optical Fiber and Cable (NZ-DSF)

Table 10-5: Abbreviated List of ITU Optical Fiber Core Standards.

The IEC includes a large collection of fiber optic standards; some examples are given in Table 10-6 [10-3].

Standard	Description
IEC 60793-1-1	Optical fibres – Part 1-1: Measurement Methods and Test Procedures – General and Guidance
IEC 61314-1:2009 (Ed. 3)	Fibre Optic Interconnecting Devices and Passive components. Fibre optic fan-outs. Generic specification
IEC 62074-1 Ed. 1	Fibre optic interconnecting devices and passive components - Fibre optic WDM devices - Part 1: Generic specification

Table 10-6: Abbreviated List of IEC Optical Fiber Standards.

10.3 Optical Connectors

This subsection considers standards that apply to optical fiber connectors. Section 4, Connectors and Interconnection Hardware, addresses the physical and mechanical description of the different types of connectors. The two main standards bodies that relate to optical connectors are TIA and IEC. The standards that they are responsible for are briefly described below.

In the US, the TIA has defined the connector standards through the EIA/TIA 604 Fiber Optic Connector Intermateability Standard (FOCIS). These standards are known as FOCIS XX or as TIA-604-XX; where each XX defines a particular connector style.

Examples of these standards for several common connector types are listed in Table 10-7 [10-6], [10-7].

TIA-604-XX	FOCIS-XX	Description
TIA-604-3-B	FOCIS 3	Fiber Optic Connector Intermateability Standard, Type SC and SC-APC
TIA-604-4-B	FOCIS 4	Fiber Optic Connector Intermateability Standard, Type FC and FC-APC
TIA-604-10A	FOCIS 10	Fiber Optic Connector Intermateability Standard, Type LC

Table 10-7: Abbreviated List of TIA Connector Standards.

Internationally the IEC has defined connector standards through the IEC 61754 standard. Each particular connector style is referred to as IEC 61754-XX. Examples of these standards for several common connector types are listed in Table 10-8 [10-8], [10-9].

IEC 61754-XX	Description
IEC 61754-4	Fibre optic connector interfaces - Part 4: Type SC connector family
IEC 61754-13	Fibre optic connector interfaces - Part 13: Type FC-PC connector
IEC 61754-20	Fibre optic connector interfaces - Part 20: Type LC connector family

Table 10-8: Abbreviated List of IEC Connector Standards.

These optical connector standards include intermateability standards that allow the same type of connectors from different vendors to mate with each other. The context of these intermateability standards, according to the TIA, is: “Intermateability standards define the minimum physical attributes of mating connector components. Fully dimensioned components are not within the scope or intent of FOCIS. The requirements of FOCIS have been selected with the objectives of ensuring that any combination of plugs and adapters conforming to the requirements of FOCIS will mechanically intermate and that intermated connector assemblies will meet their common level of performance” [10-6].

10.4 Eye Safety

Protecting the eyes from accidental laser damage is a serious issue. The danger associated with laser eye exposure must be taken seriously since we are dealing with concentrated light from either a laser source or through an optical fiber. The telecom industry typically requires long reach distances, which necessitate the use of higher powered lasers and optical amplifiers. The HPC design environment is typically characterized by its short reach distances and lower power requirements.

Eye injury is dependent on the power level, wavelength and exposure duration. The effect of power level can be made worse by the use of eye loupes, microscopes or other directly magnifying instruments. Mechanical package design, mechanical interrupts, and electrical interrupts may be used to minimize the possible eye exposure to optical energy.

Eye safety issues are covered by two sets of standards: the ANSI Z136, and the IEC 60825 standards.

The American National Standard Institute (ANSI) Z136 series is used in the US and is recognized by the Occupational Safety and Health Administration (OSHA). Listed in Table 10-9 are two of the relevant standards [10-10], [10-11], [10-12].

ANSI Z136.1-2007	American National Standard for Safe Use of Lasers
ANSI Z136.2-1997	American National Standard for Safe Use of Optical Fiber Communication Systems Utilizing Laser Diode and LED Sources

Table 10-9: ANSI Z136 Safety Standard Series.

The IEC 60825 series is recognized internationally. Listed in Table 10-10 are the two relevant standards [10-13].

IEC 60825-1	Safety of Laser Products – Part 1: Equipment Classification, Requirements and User's Guide
IEC 60825-2	Safety of Laser Products – Part 2: Safety of Optical Fibre Communication Systems (OFCS)

Table 10-10: IEC 60825 Safety Standard Series.

Both of these standards cover the 180 nm to 1 mm optical wavelength range. The wavelength, power and exposure duration are used to define the injury risk. They also address issues to avoid exposure to the optical energy.

Earlier laser classifications of Class I, II, IIa, IIIa, IIIb, and IV, which are still in use today, have been revised in ANSI Z136.1-2007. A simplified description of this set of safety standards is given in Table 10-11 [10-10], [10-11], [10-12].

Laser Classification	Simplified Description
1	The accessible laser radiation is not dangerous under reasonable conditions of use. Examples: 0.2 mW laser diode, fully enclosed 10 W Nd:YAG laser
1M	The accessible laser radiation is not hazardous, provided that no optical instruments are used, which may focus the radiation.
2	The accessible laser radiation is limited to the visible spectral range (400-700 nm) and to 1 mW accessible power. Due to the blink reflex, it is not dangerous for the eye in the case of limited exposure (up to 0.25 s). Example: some (but not all) laser pointers
2M	Same as class 2, but with the additional restriction that no optical instruments may be used. The power may be higher than 1 mW but the beam diameter in accessible areas is large enough to limit the intensity to levels which are safe for short-time exposure.
3R	The accessible radiation may be dangerous for the eye but can have at most 5 times the permissible optical power of class 2 (for visible radiation) or class 1 (for other wavelengths).
3B	The accessible radiation may be dangerous for the eye, and under special conditions, also for the skin. Diffuse radiation (as e.g. scattered from a diffuse target) should normally be harmless. Up to 500 mW is permitted in the visible spectral region. Example: 100 mW continuous-wave frequency-doubled Nd:YAG laser
4	The accessible radiation is very dangerous for the eye and for the skin. Even light from diffuse reflections may be hazardous for the eye. The radiation may cause fire or explosions. Examples: 10 W argon ion laser, 4 kW thin-disk laser in a non-encapsulated setup.

Table 10-11: International Laser Classification, With Somewhat Simplified and Approximate Descriptions. For Details, Consult the Applicable Laser Safety Standard Documents.

10.5 Wavelength Division Multiplexing (WDM)

Coarse WDM (CWDM) and dense WDM (DWDM) wavelengths, and their respective channel separations are specified by the ITU-T standards. CWDM accommodates lasers with high spectral width and/or large thermal drift. Typical wavelength separations for CWDM are in the range of 20-50 nm. Wavelength separations specified for DWDM are approximately in the range of 0.21-0.85 nm (frequency separation of 25-100 GHz). More information is contained below in Table 10-14. The smaller channel separations used for DWDM demand more precise lasers with narrow spectral width and minimum thermal drift. In contrast to CWDM, DWDM has the advantage of providing greater channel density for a given optical bandwidth allocation**.** The wide range of wavelengths used in WDM is further organized into standard bands (Table 10-12 [10-14]**)**, the names of which, in some cases, come from their historic origins. For example the "O" in "O-Band" stands for "original" since the earlier WDM systems operated near the 1300 nm wavelength. For current wavelength allocation, Table 10-13 [10-15] lists the eighteen standard CWDM center wavelengths, while Table 10-14 lists the DWDM center wavelengths for C-band for four different wavelength spacings [10-16]**.**

Allocation of Spectral Bands		
Band	**Descriptor**	**Wavelength Range (nm)**
O-band	Original	1260 to 1360
E-band	Extended	1360 to 1460
S-band	Short wavelength	1460 to 1530
C-band	Conventional	1530 to 1565
L-band	Long wavelength	1565 to 1625
U-band	Ultralong wavelength	1625 to 1675

Table 10-12: Allocation of Spectral Bands.

ITU-T G.694.2 Table (CWDM Grid)					
Channel Frequency (THz)	**Channel Wavelength (nm)**	**Channel Frequency (THz)**	**Channel Wavelength (nm)**	**Channel Frequency (THz)**	**Channel Wavelength (nm)**
235.87	1271	215.52	1391	198.41	1511
232.22	1291	212.47	1411	195.81	1531
228.67	1311	209.50	1431	193.29	1551
225.24	1331	206.61	1451	190.83	1571
221.90	1351	203.80	1471	188.43	1591
218.67	1371	201.07	1491	186.09	1611
Note: Channel spacing is defined in wavelength: delta = 20 nm (2500 GHz).					

Table 10-13: ITU-T G.694.2 Table of Standard Center Frequencies for Coarse Wavelength Division Multiplexing (CWDM Grid).

ITU-T G.694.1 Table (DWDM Grid Referenced to 193.1 THz @ 1552.52 nm)		
Channel Spacing (GHz)	**Channel Frequency (THz)**	**Channel Wavelength, λ (nm)**
12.50	f (THz) = 193.1 ± n x 0.0125	λ = c/f, where λ is the free-space wavelength in nm, f is the frequency in THz and c is the speed of light in vacuum (0.299792458e9 m/s).
25.00	f (THz) = 193.1 ± n x 0.0250	
50.00	f (THz) = 193.1 ± n x 0.0500	
100.00	f (THz) = 193.1 ± n x 0.1000	
	Where n is any integer	
Note: Channel spacing is defined in frequency (GHz).		

Table 10-14: ITU-T G.694.1 Table of Standard Center Frequencies for Dense Wavelength Division Multiplexing (DWDM Grid Referenced to 193.1 THz @ 1552.52 nm).

10.6 Package Standards

As seen in the previous sections there are standards that define essentially all aspects of the optics industry; from the ITU grid, eye safety, optical fiber and environmental concerns. The various packages incorporate these standards where appropriate but their form factors are specified by a MSA (see Subsection 10.1.2). The MSA assures that an optical package or component will have multiple supply sources allowing greater part acceptance in the market place.

Some examples of optical package MSAs are given in Table 10-15.

MSA	Website Reference for Further Details
XLMD MSA which was formed in March, 2007 to establish compatible sources of 40 Gbps Transmitter Optical Sub-Assembly (TOSA) and Receiver Optical Sub-Assembly (ROSA) for use in the 40 Gbps transceivers [40 Gbps Miniature Device (XLMD), where XL = 40]	www.xlmdmsa.org
CFP MSA which defines a hot-pluggable optical transceiver form factor to enable 40 Gbps and 100 Gbps applications, including next-generation High Speed Ethernet (40GbE and 100GbE) [C form-factor pluggable (CFP) where C = 100]	www.cfp-msa.org
300 Pin MSA for 10 Gbps and 40 Gbps Transponder formed in 2001	www.300pinmsa.org

Table 10-15: Three MSA Examples Related To Packaging Standards.

10.7 Standards Organizations Relevant to HPC Optics

A list of some of the major standards organizations is given below in Table 10-16.

Standards Organization	Description
National Electrical Code NEC http://www.nfpa.org	The NEC is a set of codes & standards defined by the National Fire Protection Association (NFPA), the document series is NFPA 7X. NFPA 76 “Standard for the Fire Protection of Telecommunications Facilities”; is probably the most applicable standard.
Electronic Industries Alliance EIA http://www.eia.org/	The EIA has its origins back to 1924; it is a trade organization representing the United States and is organized along specific product and market lines, of which the TIA is an organization, allowing each EIA sector to be responsive to its specific needs.
Telecommunications Industry Association TIA	The TIA is the leading trade association for the information, communications and entertainment technology accredited by ANSI. It represents about 600 telecommunication companies that help create universal networking and education standards that have been used worldwide.
American National Standard Institute ANSI http://www.ansi.org/	ANSI dates back to 1918 with the American Institute of Electrical Engineers (now IEEE) when it worked with four other engineering societies and three government agencies to develop a private non-profit organization. ANSI addresses standards and conformity issues to strength the US market place in the world economy. It worked with international agencies to create the International Organization for Standardization (ISO) and its sister agency the IEC to help set international goals.
International Electrotechnical Commission IEC http://www.iec.ch/	The IEC has its starts in the US in 1904 but was officially founded in 1906 in England. From the IEC website, “The IEC is the world's leading organization that prepares and publishes International Standards for all electrical, electronic and related technologies - collectively known as “electrotechnology"." It is a not-for-profit, non-governmental international standards organization.

Standards Organization	Description
International Telecommunications Union - Telecommunication Standardization Sector ITU-T http://www.itu.int/ITU-T/	The ITU has it origins back in 1865; it became a specialized agency of the United Nations in 1947. The ITU-T is one of the ITU's three sectors; representing the Telecommunication Standardization Sector. The ITU works to coordinate the shared global use of the radio spectrum, assigning satellite orbits for example; and establish world wide standards that foster seamless interconnection between the vast range of communication systems.
Occupational Safety and Health Administration OSHA http://www.osha.gov/	OSHA is an agency of the US Department of Labor created in 1970. OSHA regulations cover most private workplaces; its mission is to prevent work-related injuries, illness, and deaths.
Telcordia Technologies Telcordia http://www.telcordia.com/	Created in 1982 from what was formerly Bell Communications Research, Inc or Bellcore; based in the US, it is a telecommunications research and development company.
Multi-Source Agreement MSA Device dependent	The MSA is multi-company collaboration to develop and market a product design that they perceive as meeting a market need; an individual MSA is design dependent.
InfiniBand® Trade Association (IBTA) InfiniBand http://www.infinibandta.org/	The IBTA was founded in 1999 and its charter is to maintain and further the InfiniBand specifications. InfiniBand represents the merging of two competing designs: Future I/O and Next Generation I/O.
Open Fabric Alliance OFA http://openfabrics.org/	The OFA is an open-source project. Its goal is to develop, test, license and distribute Linux and Windows drivers and software for high-performance, low-latency networks. It supports both Ethernet and InfiniBand.

Table 10-16: List of Standards Organizations and Corresponding Websites.

10.8 References

[10-1] X40 MSA Homepage Website: http://x40msa.org/

[10-2] Welcome to XFPMSA.Org. Website: http://www.xfpmsa.org/cgi-bin/home.cgi

[10-3] The Fiber Optic Association Tech Topics Website: http://www.thefoa.org/tech/standards.htm

[10-4] Gilmore, M.: An Overview Of Singlemode Optical Fibre Specifications Website: http://www.fia-online.co.uk/pdf/Whites/wp0704.pdf

[10-5] Fiber Types according to ITU-T; Optical Communication Forum Website: http://vtsc.info/en/blog/2009/06/fiber-types-according-to-itu-t/comment-page-1/#comment-963

[10-6] The Fiber Optic Association - Tech Topics Website: www.thefoa.org/tech/focis.html

[10-7] Electronics and Telecommunications HIS Website: http://electronics.ihs.com

[10-8] Wikipedia Definition of Optical fiber connector Website: http://en.wikipedia.org/wiki/Optical_fiber_connector

[10-9] Compliance Online Website: http://www.complianceonline.com

[10-10] American National Standard for Safe Use of Lasers (ANSI Z136.1-2007) Website: http://nuigroup.com/?ACT=28&fid=16&aid=2322_Ii4NMDD5oSYnVDx TmrOr

[10-11] Laser Safety Approvals and Compliance, Valkom Laser Consulting Website: http://www3.sympatico.ca/nstoev/resource.html

[10-12] Laser Institute of America Website: http://www.laserinstitute.org/store/ANSI/112

[10-13] Encyclopedia of Laser Physics and Technology; RP Photonics Website: www.rp-photonics.com/laser_safety.html

[10-14] Kapron, F.: Some Recent Topics in Physical-Layer System Standards Website: http://www.itu.int/itudoc/itu-t/workshop/optical/s4amp01_pp7.ppt#260,4,Allocation of Spectral Bands -Sup.dsn

[10-15] G.694.2. Spectral Grids for WDM Applications: CWDM Wavelength Grid ITU Website: http://www.itu.int/rec/T-REC-G.694.2-200312-I/en

[10-16] G.694.1. Spectral Grids for WDM Applications: DWDM Frequency Grid ITU Website: http://www.itu.int/rec/T-REC-G.694.1-200206-I/en

11 Optical Computing Elements

11.1 Introduction

Optical computing, also known as all-optical processing or photonic digital processing, uses photons or beams of light to perform digital logic functions. It is hoped that all-optical computing will have higher processing speeds than electronic processors. “Although it is extremely unlikely that optical signal processing will ever be as complex and powerful as electronic processing, the increased optical functionality would have several benefits for data processing applications. Among these benefits is the ability to directly process optical headers for routing in future optical packet networks and to directly examine high-speed optical header information for optical security algorithm implementation” [11-1]. This optical fast packet switching in communications networks, including short-range interconnection networks, is feasible with small optical processors consisting of 10s to 100s of components operating at much higher speeds (100s of Gbps).

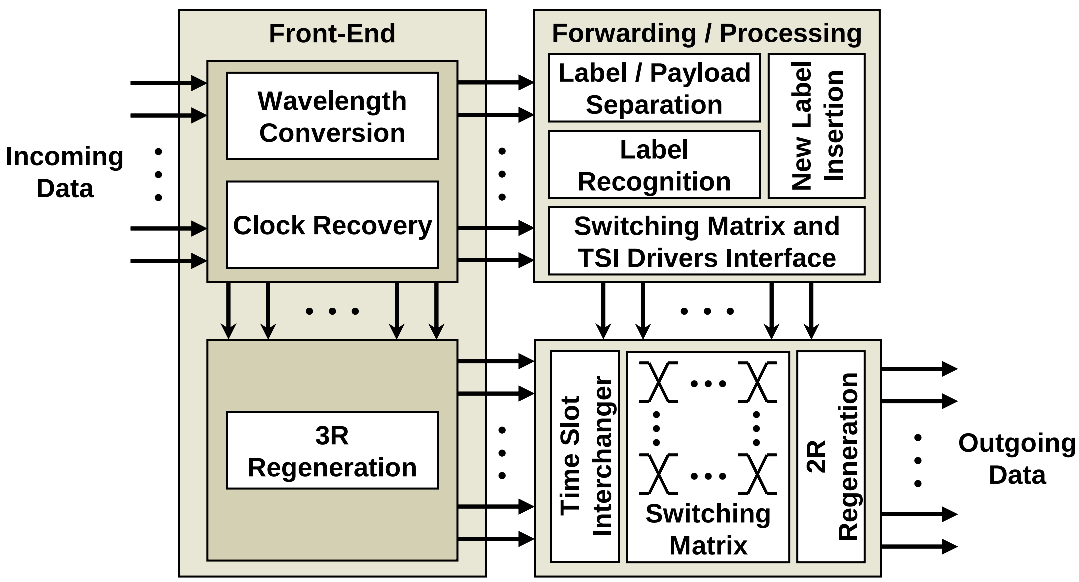

Figure 11-1: Block Diagram of an All-Optical Packet-Switched Node (From [11-2]). (40787)

Figure 11-1 shows the block diagram of a notional all-optical packet-switched node [11-2]. The required functions to realize this node include wavelength conversion, optical regeneration, optical Boolean logic and advanced optical processing circuits. Many of these functions have been demonstrated in single-gate experiments and several of the subsystems have also been demonstrated. These demonstrations have been done in a variety of technologies and have included elementary Boolean logic gates, such as AND, NOR, NAND, NOT and XOR, along with more complex functions, such as comparators, full adders and linear feedback shift registers.

In the next subsection, we will summarize the technologies currently available for optical computing and then discuss two technologies, semiconductor optical amplifier (SOA)-based devices and Self-Electro-optic Devices (SEEDs), in more detail.

11.2 Optical Computing Technologies

The underlying photonic switch technologies can be grouped into five categories based on their operating principles: electro-optic, thermal optical, optical microelectromechanical systems (MEMS), opto-optical and acousto-optic.

Electro-optic devices utilize a change in a material parameter caused by an electric field. Electro-absorption and Pockels cell modulators (see Section 3) are examples of electro-optic modulators. There are a variety of photonic switches based on electro-optic effects, including SEEDs and devices based on SOAs, lithium niobate, liquid crystals and Bragg gratings. Within the scope of our limited literature search it appears that SOA-based devices are the most common devices currently being used for all-optical processing and it appears that researchers at Sandia National Labs view SEEDs as very promising devices for all-optical logic with the potential for higher degrees of logic [11-3].

Opto-optical switches realize switching functions relying on the intensity-dependent nonlinear optic effect (which is ultrafast) in optical waveguides, such as two-photon absorption, lightwave self action (which induces the optical phenomenon of self phase modulation, SPM) and the Kerr effect (which induces the optical phenomena of four wave mixing, FWM, and cross phase modulation, XPM). They are also called optically controlled switches or all-optical switches. There are two main types: optical fiber-based switches and semiconductor-based switches. Semiconductor-based all-optical waveguide switches have many important issues to be considered before practical applications can be implemented: low operating power, ultrafast operation, high extinction ratio, room temperature operation, and polarization-independent operation [11-4].

Optical MEMS are miniature devices with simultaneous optical, electrical, and mechanical functionalities, fabricated using batch process techniques derived from microelectronic fabrication. Optical MEMS provide intrinsic characteristics for very low crosstalk, wavelength insensitivity, polarization insensitivity, and scalability. Optical MEMS-based switches are distinguished in being based on mirrors, membranes, and planar moving waveguides. The former two are free-space switches; the latter are waveguide switches. The MEMS implementations utilizing mirrors have switching speeds on the order of milliseconds which is much too slow for packet switching applications. Waveguide-based MEMS have switching speeds on the order of 100s of ns which could make them useful for packet switching applications [11-4].

Thermal optical switches are based on waveguide thermo-optic effect or thermal phenomena of materials. These devices are typically polarization-insensitive but their switching speeds are on the order of a millisecond and are too slow for packet switching applications. Switches based on waveguide thermo-optic effect are called thermo-optic switches (TOSW), which can use well-established planar lightwave circuit (PLC) technology. They are divided into two basic types: digital optical switches (DOSs) and interferometric switches. Another kind of thermo-optic switch is based on thermal effects of materials, such as thermo-capillarity optical switches, thermally generated bubble-type switches, and thermo-optic switches using coated microsphere resonators [11-4].

Acousto-optic switches are based on the acousto-optic effect in crystals such as TeO_2, in which ultrasonic waves cause periodic variations in the materials' refractive index (RI) [11-4].

The RI variations form a Bragg grating which can deflect light. This effect is used in modulators and tunable filters [11-5].

In the following sections we will briefly discuss the operating principles and demonstrated speeds of several SOA-based technologies and of SEEDs.

11.2.1 SOA-based Logic Devices

SOA-based devices can be fabricated using several different operating principles. These devices offer different functions, speeds, complexity, integratability and polarization dependence. We will briefly discuss the operating principles for several logic functions. The papers describing SOA-based logic devices generally did not list the electrical power consumption of the devices. Several of them did list the DC current in the SOAs, which is typically 300-700 mA. The voltages were not listed in these papers but from datasheets of commercially available SOAs, the typical voltage is approximately 2 V. So the power consumption in the SOAs described below is estimated to be from 0.6 W to 1.4 W.

One of the simplest operating principles for SOA-based logic devices is cross-gain modulation (XGM). It is represented schematically in Figure 11-2a. In XGM, the gain of the SOA is saturated when a strong control beam (data) is a logic high. When this occurs, the weaker continuous wave (CW) probe beam at a different wavelength will not be amplified by the SOA and the SOA output beam will be at a logic zero. When the data beam is at a logic zero, the SOA gain will recover, the CW probe beam will then be amplified and the SOA output at that wavelength will be a logic high. If a filter that passes the probe wavelength is added after the SOA, the output will be the inverse of the input data with a wavelength conversion. The speed of this device is limited by the gain recovery time of the SOA, which has been improved to allow 40 Gbps bit-by-bit wavelength conversion. A 2-bit AND function is realized by using the probe as one input and the control as the other input with the result being the following logic operation [Probe AND not(Control)].

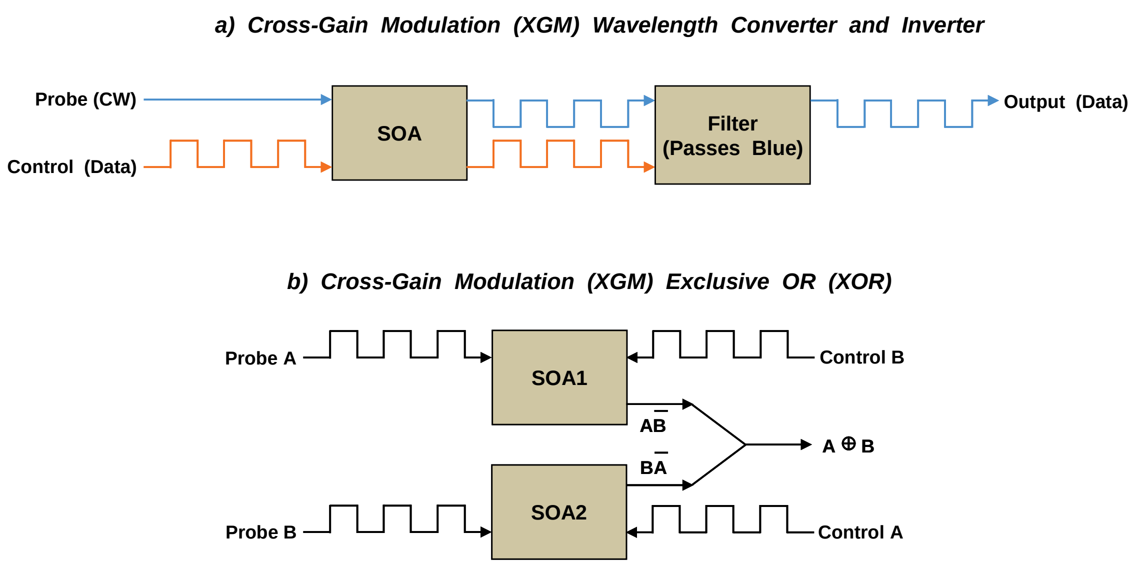

Figure 11-2: SOA Devices Utilizing Cross-Gain Modulation; (a) Wavelength Converter And Inverter, (b) Exclusive OR. (40788)

An exclusive OR (XOR) function can also be realized using XGM and two SOAs. In this configuration, the input A is the control (stronger) beam for one SOA and the probe (weaker) beam for the other, while B is the control (stronger) beam for the second SOA and the probe (weaker) beam for the first, as shown in Figure 11-2b. SOA1 performs the function $A\overline{B}$ and SOA2 performs the function $B\overline{A}$. These two functions combined perform the XOR function. In this implementation both beams are at the same wavelength so no wavelength conversion takes place. Isolators are used at each input so that the counter propagating beams do not interfere with the sources of A and B or with other components. This device has been experimentally demonstrated at 10 Gbps with no input beam polarization dependence [11-6]. A variation of this device has been demonstrated at 40 Gbps [11-7].

When a strong control beam saturates the gain of an SOA, it also changes the RI in the device which results in a phase change in light passing through it. This is known as cross-phase modulation (XPM). This phase change can also be used to construct logic devices by putting the SOA in an interferometer.

A 1x2 optical switch and wavelength converter can be made using an SOA in each arm of a Mach-Zehnder interferometer as shown in Figure 11-3. The data is introduced into one of the SOAs as the stronger control beam, while a weaker CW probe beam is split and sent to both SOAs. The control beam changes the RI in one SOA which causes a phase change in only one probe beam while the phase of the other probe beam is not changed. When the two probe beams are brought back together, if the phase difference is 0 radians, the probe beam will go to one port (output1) and if the phase difference is radians, the probe beam will go to the other port (output2), resulting in a wavelength conversion and the non-inverted data stream going to one port and the inverted data stream going to the other. This device requires less optical input power and offers a higher extinction ratio than XGM-based devices. However, just like XGM-based devices, its speed is limited by carrier recovery time in the SOAs [11-1]. This limitation can be overcome by using more complex structures.

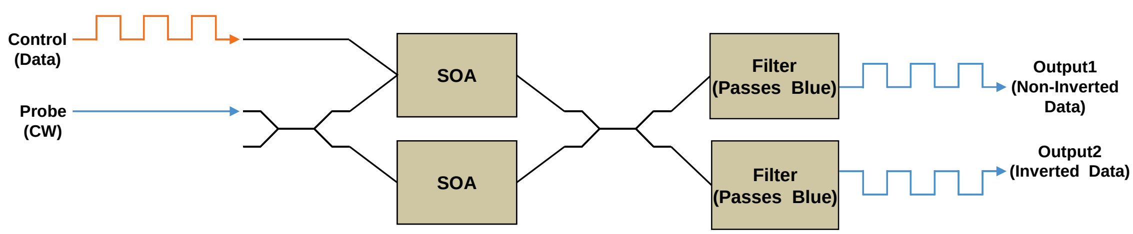

Figure 11-3: SOA Device Utilizing Cross-Phase Modulation To Implement A 1x2 Switch And Wavelength Converter. (40789)

An example of a more complicated circuit that has been demonstrated is a 2-bit comparator operating at 10 Gbps [11-7]. In routers used in a packet switched network, when two packets are simultaneously destined for the same output port, the packet with the highest priority is sent to that port first. A photonic digital circuit able to compare two multi-bit Boolean numbers is needed for an optical router and the optical 2-bit comparator presented here is an early demonstration of such a circuit. The upper part of Figure 11-4 shows the block diagram for the comparator and the lower part shows the basic hardware implementation.

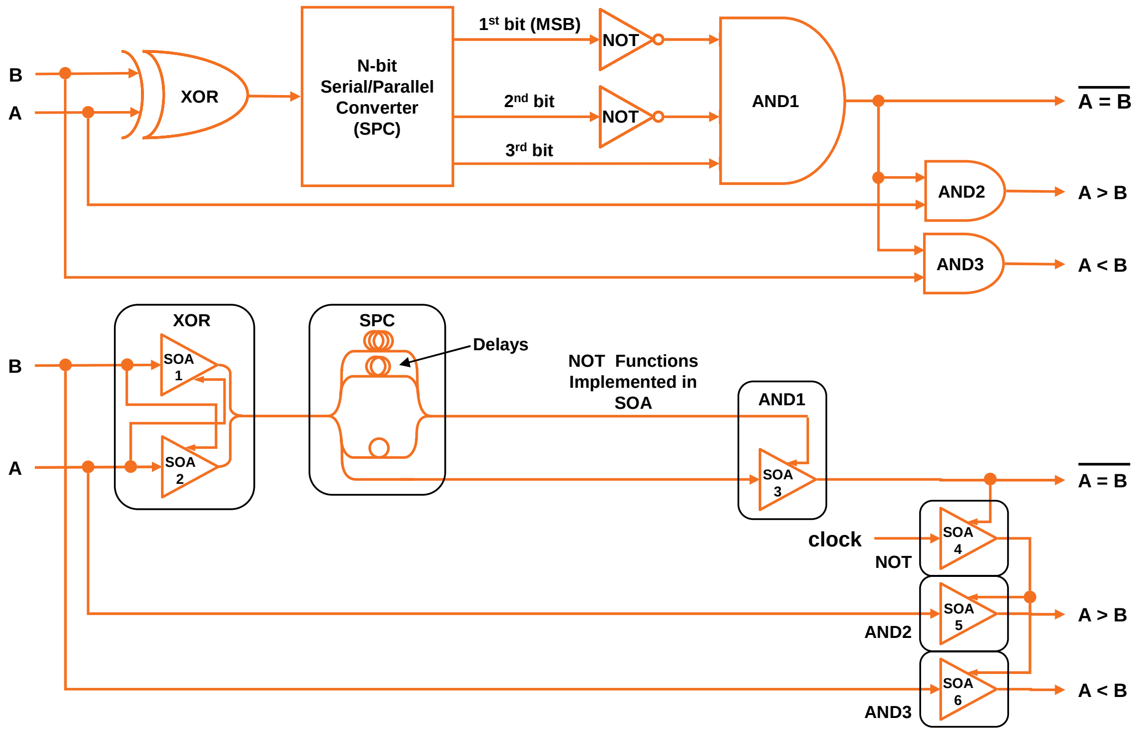

Figure 11-4: Block Diagram (Upper) And Basic Hardware Implementation (Lower) Of An N-Bit Photonic Comparator (From [11-8]). (40790)

The N-bit Boolean numbers are input as serial bit streams at A and B with the most significant bit (MSB) arriving first. The first operation is a bit-wise XOR using the circuit previously presented in Figure 11-2b. Next the serial bit stream is parallelized in a special serial-to-parallel converter (SPC). This SPC does not wait for all of the bits to arrive before sending bits out and the early arriving bits are sent out on different lines with different amounts of delay. As shown in Table 11-1, in the first bit period, the first bit is output (undelayed) on the first line. In the second bit period, the first bit is output (1-bit period delayed) on the second output line and the second bit is output (undelayed) on the first output line. In the third bit period, the first bit is output (2-bit period delayed) on the third output line, the second bit on the second output line (1-bit period delayed) and the third bit on the first output line (undelayed). This continues until all of the bits have arrived. The least significant bit (LSB) is the last to arrive and is sent out on the first output line only. Zeroes are sent over the output lines that the data has not arrived at yet, e.g., when the first bit is sent on the first output line, all of the other lines have zeroes. A guard band of N-bit times is required between numbers so the SPC can reset. Next the delayed bits are inverted and all of the bits are ANDed to give the output inverse of A=B. The AND function is implemented using a circuit similar to the one presented in Figure 11-2a. This data is ANDed with A and B to give the outputs A>B and A<B, respectively. The outputs are N-bits long with all zeroes if the output is false and a single '1' with the rest of the bits being zero if the output is true. The comparator can be implemented using six SOAs regardless of the length (N) of the numbers being compared. This demonstration was built with individual components connected with fibers. The SOAs are approximately 2 mm in length and require roughly 5 dBm for the pump signal and -9 dBm for the probe signal. The electrical power consumption was not given in the paper.

	Bit Period 4	Bit Period 3	Bit Period 2	Bit Period 1
Output Line 4	Bit 1	0	0	0
Output Line 3	Bit 2	Bit 1	0	0
Output Line 2	Bit 3	Bit 2	Bit 1	0
Output Line 1	Bit 4	Bit 3	Bit 2	Bit 1

Table 11-1: Serial-To-Parallel Converter Outputs At Different Times.

Another circuit that has been used in several demonstrations is the SOA-based ultrafast nonlinear interferometer (UNI). The UNI gate can be configured as an AND or an XOR which have been demonstrated at 100 Gbps [11-9] and 85 Gbps [11-10], respectively. UNI gates have also been used in demonstrations of binary adders [11-11], data encoders/decoders and other circuits.

The UNI gate relies on a phase shift between two orthogonally polarized components of a clock input pulse. The block diagram of an XOR based on a UNI gate is shown in Figure 11-5. The clock input pulse arrives from the left and is split into its polarization components using cross-spliced polarization maintaining fiber (PMF). One of the polarization components is delayed by the PMF before they both enter the SOA as probe signals. The data input pulses A and B are the control signals entering the SOA from the right. The presence of a control pulse causes a phase shift in the probe polarization component that it is coincident with the control pulse. After the SOA, a section of PMF delays one polarization component so the two are again coincident and another cross-spliced PMF recombines the two components. If only one of the controls is present, the differential phase shift between the probes results in polarization rotation, hence inducing constructive interference, and the output will be 1. If both controls are present or absent, the phase change in the probes is adjusted equally to keep the interferometer in balance, and the output will be 0 [11-12]. This gate requires very low energy per operation (<5 fJ compared with approximately 200 fJ in a similar demonstration without an SOA) but it also requires rigid synchronization of the light pulses thus making it vulnerable to timing jitter.

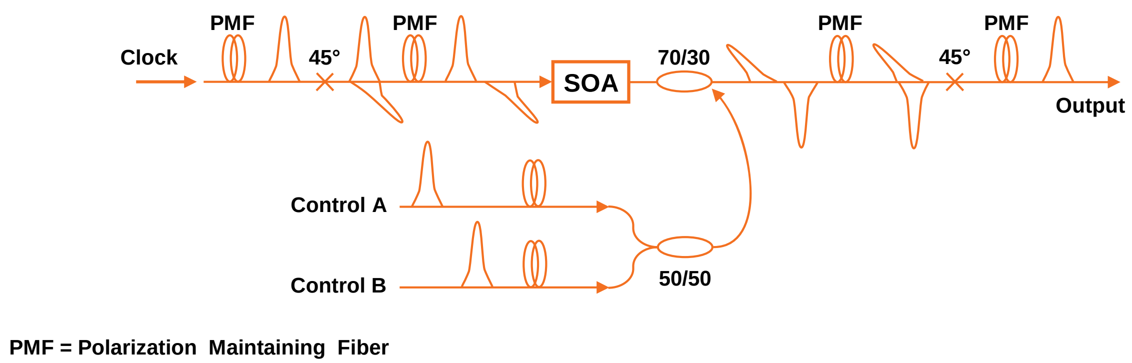

Figure 11-5: Block Diagram of an XOR Based On an Ultrafast Nonlinear Interferometer (UNI) Gate (From [11-12]). (40791)

11.2.2 SEEDs

Self-Electro-optic Devices (SEEDs) [11-13] and a variant on the technology, Symmetric SEEDs (S-SEEDs) [11-14], were demonstrated in the late 1980s by workers at Bell Labs. Since then, significant research has been carried out at Bell Labs, and more recently at Sandia National Labs, to develop them into a viable optical computing technology. They have the potential for high speed and, with micro-optic interconnect technology, they also have the potential for dense device integration. The main drawbacks of the technology are that they require high-power clock sources and accurately timed reset and clock pulses that make them vulnerable to control signal distribution issues causing timing jitter. Individual S-SEEDs can be configured for NAND/AND or NOR/OR operations and multiple devices can be connected together to construct more complicated circuits. One example is an XOR function that was implemented using three S-SEEDs [11-3]. While the integrated S-SEEDs along with the micro-optical interconnect are quite small (~3 mm x 1 mm), at this time they still require a rack of equipment to generate and control the optical pulses. Another example of a more complex circuit is a shift register proposed in [11-15].

SEEDs are vertical P-I-N devices that resemble VCSELs with a DBR mirror on the bottom side of the device only and not on the top side as shown in Figure 11-6A [11-16]. To construct an S-SEED, two of these structures are connected in series and a large diode, acting as a capacitor, is connected in parallel with them as shown in Figure 11-6B. At wavelengths near the MQW band gap, the optical absorption, and hence the device's photocurrent response, is strongly dependent on applied bias via the quantum-confined Stark effect (QCSE). Figure 11-6C shows this effect for a device with its absorption edge near 1550 nm. The resulting photocurrent versus bias curves for an S-SEED are shown in Figure 11-6D. These curves intersect at three points, producing optical bistability as a function of input power. With balanced inputs (either with light or with no light), the device is in one of these two stable states. Unbalanced inputs cause it to switch states.

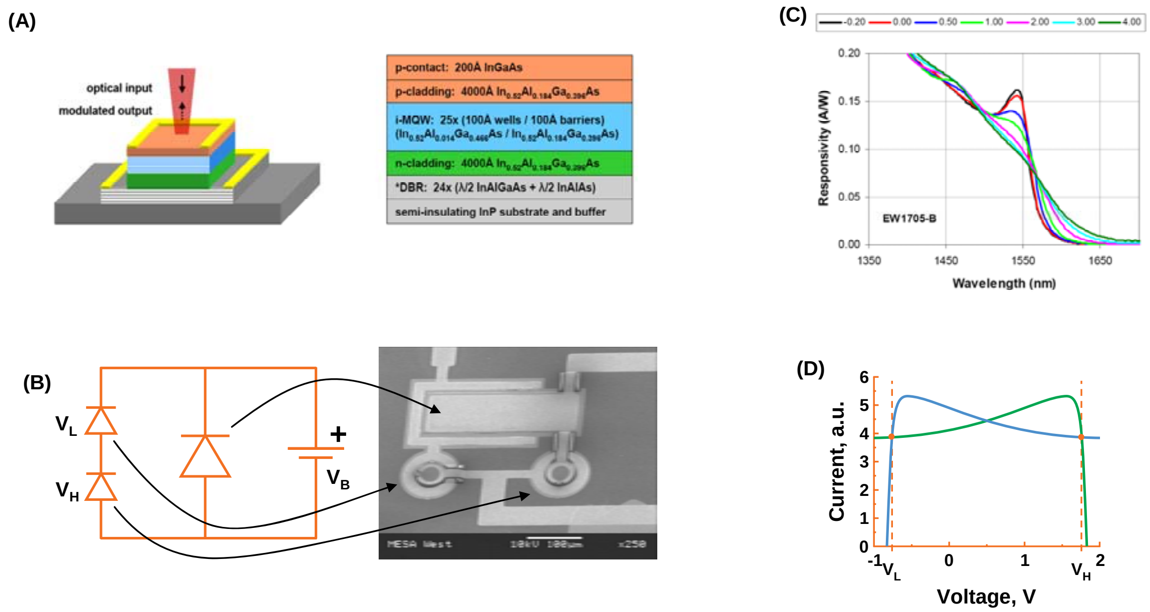

Figure 11-6: (A) Diagram Of Individual SEED (Left) And Typical Layer Stack Up (Right); (B) Schematic Diagram of S-SEED With Arrows Indicating Individual Components On the Photomicrograph of A Device; (C) SEED Responsivity Versus Wavelength; and (D) Photocurrent Versus Voltage Plot of An S-SEED (From [11-16]). (40812)

11.3 Conclusions

It appears that the majority of the research in optical computing/processing is geared toward all-optical networks and all-optical packet processing for the commercial telecom market. "Some of the functions in today's commercial (telecom) systems such as optical amplification and add/drop multiplexing are already performed in the optical domain using either Erbium Doped Fiber Amplifiers (EDFAs) or Raman Amplifiers, in the first case, and ROADMs [reconfigurable optical add/drop muxes] in the second one [11-17]." However, other components necessary for all-optical networks and all-optical packet processing are still in the research stage. While significant progress is being made, these components will, likely, not be available for commercial systems for 5-10 years.

11.4 References

[11-1] Poustie, A.: Semiconductor Optical Amplifiers Light Up All-Optical Signal Processing. Photonics Spectra, 8:62-66, 2007.

[11-2] Raffaelli, C., et al.: Photonics in Switching: Architectures, Systems and Enabling Technologies. Computer Networks, 52:1873-1890, 2008.

[11-3] Keeler, G. A., et al.: All-Optical XOR using 1550-nm High-Speed S-SEEDs and Integrated Micro-Optics. Conference on Photonics in Switching. PS 2008, pp. 1-2, 4-7 August 2008.

[11-4] Ma, X., and G-S Kuo: Optical Switching Technology Comparison: Optical MEMs Vs. Other Technologies. IEEE Communications Magazine, 41(11):S16-S23 (November) 2003.

[11-5] Dutton, H. J. R.: Understanding Optical Communications. IBM Redbook 1998, http://www.redbooks.ibm.com/redbooks/pdfs/sg245230.pdf.

[11-6] Kim, J. H., et al.: All-Optical XOR Gate Using Semiconductor Optical Amplifiers Without Additional Input Beam. IEEE Photonics Technology Letters, 14(10):436-38 (October) 2002.

[11-7] Webb, R. P., et al.: 40 Gbit/s All-Optical XOR Gate Based On Hybrid-Integrated Mach-Zehnder Interferometer. Electronics Letters 39(1):79-81 (January) 2003.

[11-8] Scaffardi, M., et al.: Photonic Processing for Digital Comparison and Full Addition Based on Semiconductor Optical Amplifiers. IEEE Journal of Selected Topics in Quantum Electronics, 14(3):826-833 (May/June) 2008.

[11-9] Hall, K. L., and K. A. Rauschenbach: 100 Gbit/s Bitwise Logic. Optics Letters, 23:1271-1273 (August) 1998.

[11-10] Bintjas, C., et al.: Ultrafast Nonlinear Interferometer (UNI)-Based Digital Optical Circuits and Their Use In Packet Switching. Journal of Lightwave Technology, 21(11):2629-37 (November) 2003.

[11-11] Poustie, A .J., et al.: All-Optical Full Adder With Bit-Differential Delay. Optical Communications, 168:89-93 (June) 1999.

[11-12] Zhang, M., et al.: All-Optical XOR Logic Gates: Technologies and Experiment Demonstrations. IEEE Optical Communications, 43(5):S19-S24 (May) 2005.

[11-13] Miller, D. A. B., et al.: Novel Hybrid Optically Bistable Switch: The Quantum Well Self-Electrooptic Effect Device. Applied Physics Letters, 45:13-15 (July) 1984.

[11-14] Lentine, A. L., et al.: Symmetric Self-Electro-Optic Device: Optical Set-Reset Latch. Applied Physics Letters, 52:1419-21 (April) 1988.

[11-15] Tian, B., W. van Etten, and W Beuwer: Ultrafast All-Optical Shift Register And Its Perspective Application For Optical Fast Packet Switching. IEEE Journal of Selected Topics in Quantum Electronics, 8:722-728, 2002.

[11-16] Keeler, G. A., et al.: High-Speed Switching of a 1.55-μm Symmetric SEED. Conference on Lasers and Electro-Optics, 2008 and 2008 Conference on Quantum Electronics and Laser Science. CLEO/QELS 2008, pp.1-2, May 4-9, 2008.

[11-17] Monteiro, P., et al.: Challenges and Trends in Optical Networking: A Bottom-Up Approach. 5th International Conference on Broadband Communications, Networks and Systems. BROADNETS 2008, pp. 460-469, September 8-11, 2008.

12 Telecommunication Industry Trends

This section provides some insight into current technical trends in the telecommunications industry. The majority of the discussion refers to long-haul dense wavelength division multiplexing (DWDM) systems and is organized around four basic topics: photonic integration; advanced modulation formats; multiplexing approaches; and autonomous re-configurability. The final discussion refers to current trends in wideband radio-on-fiber (RoF) telecommunications.

12.1 Photonic Integration

Whether it is at telecommunication seminars, conferences, trade journal advertisements, or the technical literature, the participants and organizers of these venues at some point or other emphasize the rapidly increasing global bandwidth requirements for data traffic. Often cited is the recently observed bandwidth growth for just the web alone, which is on the order of 75% per year. Some refer to this growth as the bandwidth equivalent of Moore's law. For the most part this growth is driven by the internet and by the proliferation of portable and hand held multi-media communication devices. For example, today's internet traffic in the US alone exceeds 5 Tbits of data every second, according to AT&T. By 2015, it is projected to be 1000 times greater [12-1]. If they are to stay in business, the various service providers (SPs) must respond to this demand under the constraints of a modestly growing customer base with corresponding modest revenue growth. On the other hand, the cost of optical components and subassemblies using discrete and hybrid components is projected to grow faster than revenues [12-2]. For the SPs to be economically viable, component costs must come down, and monolithic photonic integration is seen as a way to accomplish this. Figure 12-1, taken from Reference [12-2], depicts this largely market-driven trend.

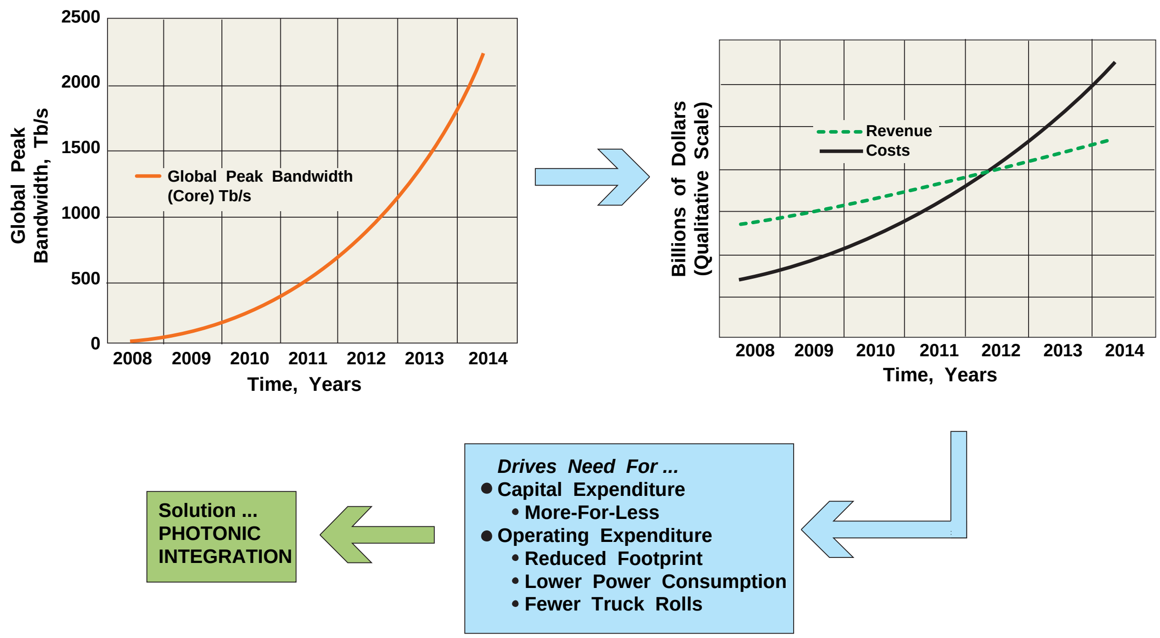

Figure 12-1: Cost of Optical Hybrid Components Is Outpacing Service Provider Revenue Growth Due to Escalating Bandwidth Demands: A Motivation for Photonic Integration. (40452)

Photonic integrated circuits (PICs) can be fabricated on III-V materials or silicon. Most commercial PICs for DWDM are fabricated on an InP substrate. For example, a PIC for DWDM transmitter applications might include several edge-emitting lasers, modulators, switches, and array waveguide gratings (AWG), all integrated on a single chip. Such a photonic integration strategy has been successfully adopted by several transceiver developers for long-haul systems, including Lucent/Bell Labs and a fairly new company, Infinera, who is today the leading maker of InP-based photonic integrated circuits (PICs) for long-haul DWDM systems Figure 12-1. Figure 12-2 shows an example of an Infinera InP transmitter chip that integrates around 50 optical components onto a single die. A research group at the University of California, Santa Barbara, has demonstrated an optical router that integrates 200 “functional elements” on a 4.25×14.5 mm InP chip and operates at 40 Gbps [12-3]. Projections are that PICs of the future (next 5 to 10 years) might incorporate several hundreds if not a thousand optical components all integrated monolithically on a single chip [12-4].

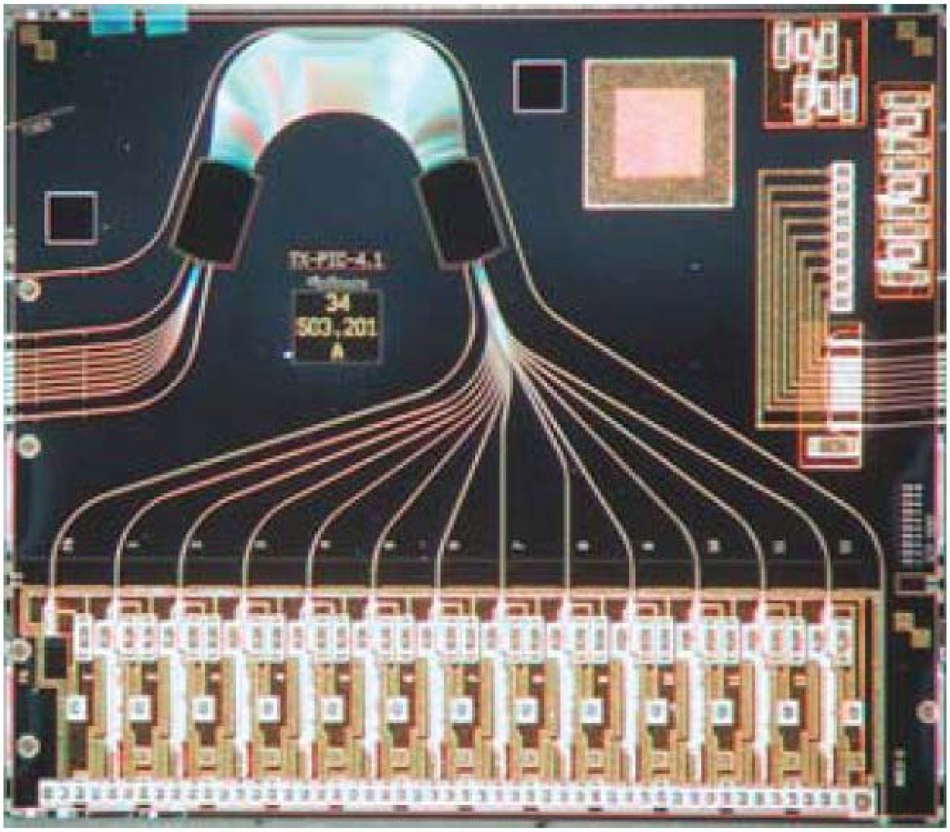

Figure 12-2: Photonic Integrated Laser Transmitter Array on InP from Infinera (About 50 Optical Components Are Integrated On Chip). (40458)

Another example of photonic integration has been successfully demonstrated by Luxtera in their active cable products. Luxtera's approach uses a single 1550 nm, group III-V edge emitting laser mounted directly onto a silicon CMOS photonic chip. On the chip the continuous wave (CW) light from the laser is coupled into optical waveguides that split into four paths that route the light to four separate monolithically integrated Mach-Zehnder (MZ) optical modulators, also fabricated on the silicon die. All electronic circuitry is implemented in CMOS including the driver for the MZ modulators. Four separate modulated channels are therefore possible with each channel capable of modulation rates up to 10 Gbps for an aggregate transmit capability of 40 Gbps.

A third photonic integration approach, although as yet at the research stage, is being pursued by researchers at Intel and the University of California, Santa Barbara. In this approach an InP-based wafer containing the lasing media (InGaAlAs) is bonded at relatively low temperature to a pre-patterned silicon-on-insulator (SOI) wafer which contains optical waveguides for guiding the light as well as CMOS circuitry for driving active components [12-5]. The InP substrate is eventually etched away leaving behind just a thin membrane of the active III-V lasing media. The emanating light from the active region is evanescently coupled to the silicon waveguides beneath the III-V material. Effectively, a hybrid, "silicon" laser is thus fabricated. The goal is to achieve this bonding at wafer level for large volume production and take advantage of the best of both worlds – planar lightwave circuits (PLC) and mature CMOS circuitry integrated on silicon, together with InP-based materials optimized for active optical functions including laser transmitters.

An interesting consequence of photonic integration is, by necessity, component miniaturization and an increase in on-chip power dissipation density. In fact, to maintain device temperatures consistent with long term reliability, overall efficiencies of devices and supporting electronic circuitry must increase and be commensurate with the number of active optical components included on a chip.

12.2 Advanced Modulation Schemes

For decades the modulation format for optical telecommunication links has been binary ON–OFF-keying, non-return-to-zero (OOK-NRZ). In directly modulated schemes, for example, the laser current is driven to obtain maximum light intensity for the “1” bit and then reduced to a much lower level for the “0” bit (the laser current is not completely shut off and is maintained above the threshold current for all bits to avoid turn-on delay and reduced frequency response). At first the demand for greater bandwidth and higher aggregate data rates was met with wavelength-division multiplexing for long-haul systems. However, as the demand for greater bandwidth and higher data rates continues to grow, long-haul DWDM systems especially, are implementing more advanced modulation schemes to increase optical link spectral efficiency and reach (spectral efficiency is defined as the ratio of the information bit rate to the total bandwidth consumed [12-6]). The advanced schemes began appearing in the first half of the present (2000-2010) decade. Initially OOK modulation was augmented with optical phase modulation in a non-information-bearing way to mitigate chromatic dispersion and other fiber-related degrading effects [12-6]. These formats include duobinary, alternate mark inversion (AMI), chirped return-to-zero (CRZ), and alternating-phase (AP) OOK formats such as carrier-suppressed return-to-zero (CSRZ) [12-7] (discussed in detail in Section 1). More recently optical phase-shift-keyed (PSK) modulation has been used to carry the information in the optical phase itself and to extend the reach of DWDM long-haul systems even further. Because there is no phase reference available in direct-detection receivers, the phase of the preceding bit is used as a relative phase reference for demodulation, resulting in differential-phase-shift-keying (DPSK). In DPSK the information is carried in the optical *phase changes* between bits [12-6] (see also Subsection 1.4.1 in Section 1).

Table 1 shows a comparison of optical spectral efficiencies for some of the advanced modulation formats currently being implemented or being considered for future long haul telecommunication systems [12-8]. This table was generated by the authors of Reference [12-8] based on their studies and simulations of data communication links for DWDM systems. One of their conclusions is that although 160 Gbps was included in their study, it seems that for all intents and purposes such fast rates are impractical (cost of implementation) for systems in the foreseeable future. On the other hand, differential quadrature phase shift keying (DQPSK) at 40 Gbps has already been demonstrated in laboratory settings using InP PICs [12-9] and new products are currently under development that soon will be commercially available (see, for example, Mintera and Opnext websites;[12-10], and [12-11]). Even more sophisticated modulation schemes have been recently demonstrated by researchers using sophisticated 128 QAM at 14 Gb/s data speed [12-12].

Example of Modulation Format	*Binary OOK*			*DPSK or Duobinary*		*4-Level (DQPSK)*		*16-Level*	
Spectral Efficiency (b/s/Hz)	0.2	0.4		0.8		1.6		3.2	
Bit-Rate per λ (Gb/s)	10	10	40	10	40	40	160	40	160
Inter-λ Spacing (GHz)	50	25	100	12.5	50	25	100	12.5	50
Number of λ's in C-Band	80	160	40	320	80	160	40	320	80
C-Band Capacity (Tb/s)	0.8	1.6		3.2		6.4		12.8	

Figure 12-3: Comparison of Spectral Efficiencies for Advanced Modulation Formats Relative to Binary OOK for DWDM-Based Optical Links and Data Rate Capacity in C-Band for Several Choices of Bit Rates (C-Band is A Standard Telecommunications Wavelength Band Extending From 1530 nm to 1565 nm; 100 GHz Frequency Spacing is Equivalent to 0.80 nm Wavelength Spacing at Mid-C-Band). (40460)

12.3 Multiplexing Approaches

Employing various multiplexing schemes is another way to increase the aggregate bandwidth and “pack” more bits onto the fiber. Three general approaches are possible based on time division, wavelength division, and space division multiplexing. All three approaches have, and will be major considerations for future telecommunication networks. Briefly, the approach and trends for increasing fiber capacity using these three multiplexing schemes are described below.

12.3.1 Time Division Multiplexing (TDM)

Conceptually TDM is the most straightforward approach to increase channel capacity. Since TDM is based on time slots, higher symbol (Baud) rates essentially means more bits can be packed into a given time slot for increased capacity. However, technologically it is quite challenging to increase rates significantly above 10 Gbps. For example, most deployed VCSEL-based systems are today limited to 10 Gigabaud (GBd) (10 Gbps for OOK), primarily because of current density limitations and reliability considerations associated with the conventional VCSEL structure. Nevertheless, research work is on-going for development of advanced VCSEL structures that potentially will be capable of significantly higher modulation rates [12-13], [12-14]. Even so, it is anticipated that VCSELs for practical systems will be limited to no more than 20 Gbps modulation rates, at least for the foreseeable future. For modulation rates around 20 Gbps DFB lasers can be directly modulated. Beyond 20 Gbps DFB lasers with external modulators are used. For most deployed systems today, however, the symbol rates remain at or below 10GBd.

12.3.2 Wavelength Division Multiplexing (WDM)

WDM is a way to increase capacity by assigning different data streams, or communication channels to different wavelengths (optical carriers) that coexist and propagate on the same optical waveguide. In the long-haul telecommunications industry dense wave division multiplexing (DWDM) schemes are ubiquitous. Many deployed systems work on the so-called ITU grid of 100 GHz frequency spacing (0.8 nm wavelength spacing at 1550 nm), but the trend is to develop systems with 50 GHz, and even 25 GHz, spacing. The tighter wavelength spacing requires extremely tight temperature control for the laser transmitters (using thermoelectric coolers (TECs)) to achieve the required wavelength stability (e.g., less than 0.1 degree C). For shorter reaches (less than about 10 km) coarse wave division multiplexing (CWDM) is also being considered [12-15]. Here the wavelength spacing is standardized to 20 nm. In contrast to DWDM systems, CWDM-based laser transmitters can run without TECs resulting in significant power savings compared to DWDM transmitters, but at reduced aggregate bandwidth.

12.3.3 Space Division Multiplexing

In space division multiplexing capacity is increased by having multiple parallel physical channels. The objective of this approach is to increase the packing density of parallel optical waveguides (also referred to as "parallel optics") to achieve large aggregate bandwidths. It is likewise conceptually very straightforward, yet technologically difficult to implement in many cases. A first level implementation of parallel optics is to bundle several optical fibers in a single ribbon. Many active cable products employ this type of interconnect between transceivers. Finisar, for example, is currently shipping 40 Gbps transceivers known as their "Quadwire" active cable. It incorporates an 8-fiber ribbon cable that provides a 40 Gbps aggregate link via four 10 Gbps full-duplex channels at 850 nm. The corresponding transceivers use linear VCSEL and photodetector array pairs that couple to their respective fiber.

850 nm VCSELs are the workhorse for short-haul LANs. Generally the industry trend is to monolithically integrate larger arrays of VCSELs and interface these with multiple lanes of optical waveguides to achieve even larger aggregate bandwidths. IBM's Terabus program, for example, demonstrated a 4 x 4 VCSEL (photodetector) array coupled to an FR4-based "optoboard" with embedded polymer optical waveguides on a 62.5 micron pitch. In this case a complete operational high speed data link incorporating 16 bidirectional channels, with each channel supporting up to 15 Gbps data rate was realized between a pair of transceivers separated by approximately 15 cm of the FR4 optoboard [12-16]. Another good example of space division multiplexing is the InP edge-emitting, DFB laser transmitter array (10 x 10 Gbps) shown in Figure 12.2 and described in detail in [12-17].

12.4 Reconfigurable Networks

A critically important consideration in the long-haul telecommunications industry is the ability to set up the routing of data traffic through the various network access points (optical nodes) of a DWDM-based system. A key component that enables this functionality is the optical add/drop multiplexer (OADM) as shown conceptually in Figure 12.4. Up until a few years ago these devices could route 1 to 4 channels (wavelengths) at the optical nodes, but these were fixed, that is, a client port was fixed in assignment and in direction. To change the network configuration, a patch panel or optical switch had to be used. Around the 2004-2005 time frame reconfigurable OADMs, or ROADMs, were introduced that incorporated remotely reconfigurable optical switches which allowed all 40 channels of a 100 GHz-spacing C-band DWDM system to be accessed at an optical node (see Figure 1-29 in Subsection 1.7 for a more detailed OADM/ROADM schematic). The result was more agility in network re-configuration and reduced operating costs for the service providers (actual revenues increased significantly as a result of ROADM deployment – comment by Cisco Systems engineer at ROADM workshop, 2009 OSA/OFC/NFOEC). Today even more flexible, multi-degree and omni-directional ROADM devices are being introduced into DWDM networks to further reduce service providers' operating costs. These devices have several monitoring capabilities as well. Ultimately the goal is to have the network be able to reconfigure and optimize routing of data traffic autonomously.

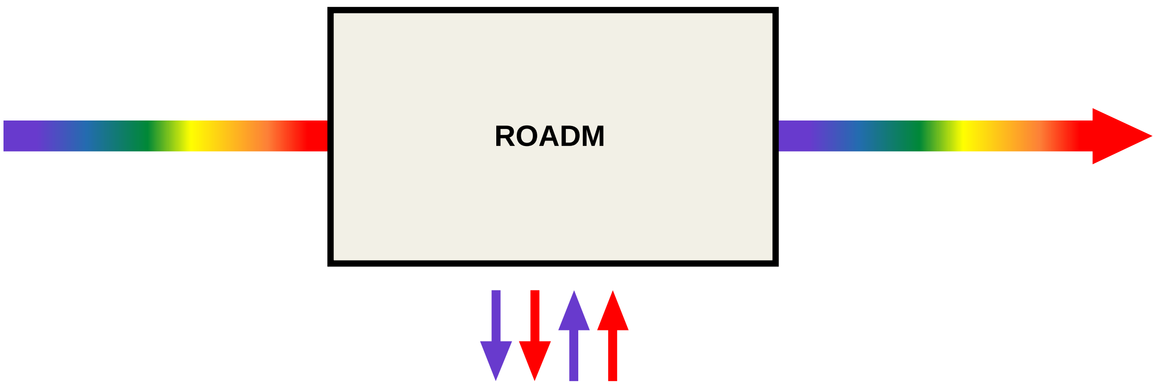

Figure 12-4: Conceptual Description of ROADM = Reconfigurable Optical Add/Drop Multiplexer (An Optical Subsystem Capable Of Selective And Automatic Removal Or Addition Of Individual Wavelengths From An Optical Fiber). (40594)

12.5 Radio-Over-Fiber (RoF) Links

There is growing interest in RoF systems that provide access to ultra-broadband high frequency radio systems (e.g., IPTV, HDTV, or Wireless HD applications) at multiple access points in public establishments such as airports, malls, train stations, conference centers, and eventually home and small business offices [12-18]. These systems are targeting the 57-64 GHz license-free frequency band for multi-Gbps wireless communication. These picocell access points broadcast a 60 GHz carrier for wireless reception by the end user. However, because of the very limited range of the 60 GHz carrier (less than 10 m), the "base stations" need another "carrier," to interconnect the network. Fiber-optics provides the solution, i.e., the 60 GHz wideband signals are used to modulate a laser for transmission of the information along a fiber network. Furthermore, there is on-going research to consider including other wireless services such as the existing Wi-Fi (2.4-GHz) and WiMAX (5.8-GHz) systems and to modulate the 60 GHz carrier with this information as well [12-19].

12.6 References

[12-1] Oyo, S. J. B., L. Sebastian, and W. T. Tsang: Planar Re-Growth Simplifies Photonic Integration of InP Chips. Compound Semiconductors, pp. 14-16 (March) 2009.

[12-2] Smith, D.: Scalable Photonic Integration: the Hybrid Solution. Presentation at Optoelectronics Industry Development Association (OIDA), October 2008, Monterey, CA.

[12-3] Optics.org web site: On-Chip Optics: The Route to Success. Work of Steve Nicholes and Milan Mašanović at UCSB, April 24, 2009. http://optics.org/cws/article/research/38791

[12-4] Smit, Meint: Remarks made at the OIDA Conference, October 2008, Monterey CA.

[12-5] Fang, A. W., H. Park, O. Cohen, R. Jones, M. Paniccia, and J. Bowers: Electrically Pumped Hybrid AlGaInAs-Silicon Evanescent Laser. OPTICS EXPRESS 14(20):9203-9210 (October) 2006.

[12-6] Couch, L. W.: Digital and Analog Communication Systems, 5th Edition. Prentice Hall, 1997.

[12-7] Gnauck, A. H., and P. J. Winzer: Optical Phase-Shift-Keyed Transmission. IEEE Journal of Lightwave Technology, 23(1):115-130 (January) 2005.

[12-8] Saleh, A., A. M., and J. M. Simmons: Evolution Toward the Next-Generation Core Optical Network. IEEE Journal of Lightwave Technology, 24(9):3303-3321 (September) 2006.

[12-9] Infinera web site: http://www.infinera.com/technology/pic_roadmap.html

[12-10] Mintera web site: http://www.mintera.com/manswerpages/mi40000xs.htm

[12-11] Opnext web site: http://www.opnext.com/products/pdf/OEB-NOV2009.pdf

[12-12] Nakazawa, M.: Optical Quadrature Amplitude Modulation (QAM) With Coherent Detection Up to 128 States. 2009 OSA/OFC/NFOEC, paper No. OTHG1, March 24-26, 2009, San Diego, CA.

[12-13] Chen, C., and K. C. Choquette: Push-Pull Modulation of Coupled-Cavity VCSELs. 2009 OSA/OFC/NFOEC, paper No. OTuK3, March 24-26, 2009, San Diego, CA.

[12-14] Ledentsov, N.: Novel Concepts for High Speed VCSELs. 2009 OSA/OFC/NFOEC, paper No. OTuK6, March 24-26, 2009, San Diego, CA.

[12-15] Rochus, S.: Photonic Integrated Circuits for 100GbE. OIDA Workshop100Gb Ethernet: The Next Challenge for Communications Systems, August 29, 2006, San Jose, CA.

[12-16] Doany, F. E., et al.: Terabus: A 160-Gb/s Bidirectional Board-Level Optical Data Bus. Proceedings of the 20th Annual Meeting of the IEEE Lasers and Electro-Optics Society, pp. 545-546, October 21-25, 2007, Lake Buena Vista, FL.

[12-17] Nagarajan, R.: Large-Scale Photonic Integrated Circuits. IEEE Journal of Selected Topics in Quantum Electronics, 11(1):50-65 (January/February) 2005.

[12-18] Waterhouse, R. B., and D. Novak: Broadband Antennas and Efficient Electro-Optic Interfaces for Fiber-Radio Applications. 2009 OSA/OFC/NFOEC, paper No. OTuK6, March 24-26, 2009.

[12-19] Hsueh, Y-T., et al.: Generation and Transport of Independent 2.4 GHz (Wi-Fi), 5.8 GHz (WiMAX), and 60-GHz Optical Millimeter-Wave Signals on a Single Wavelength for Converged Wireless over Fiber Access Networks. 2009 OSA/OFC/NFOEC, paper No. OTuJ1, March 24-26, 2009.

13 Implementation Roadmap

13.1 Introduction

As bandwidth and data rates continue to increase in high performance compute (HPC) systems, it is likely that within the next several years portions of the electrical interconnect system will fall short of the performance requirements for such systems. Gradually, electrical interconnects will give way to their optical counterparts especially for the longer reaches where the electrical bandwidth-distance product is exceeded. Just as active optical cables are finding application in data centers for rack-to-rack interconnects, so too will variations and shorter versions of such cables be adapted to intra-board and back-plane interconnects. Among the key technical advances that will determine how quickly this evolution will take place are the reduction in optical transceiver power consumption and packaging miniaturization. As other sections in this Reference Guide have indicated, the extensive use of optical interconnects for on-board interconnects will demand a reduction in the energy-per-bit metric by at least an order of magnitude from the values that are typical of active optical cables in use today. Similarly, packaging of the entire transceiver must be miniaturized and the interface to silicon-based high speed SerDes I/O must be optimized for high speed and high fidelity data transfer.

To facilitate the introduction of optical interconnects into specific government HPC systems, this section describes a chronological list of suggested actions or experiments that would be initiated when the specific optical technology is made available (either commercially or through special arrangement of loaned experimental parts). Therefore, some of the items listed have two phases. The goal of the first phase is to determine whether a particular technique or emerging technology will be commercially viable. If the answer is yes, then the second phase is envisioned as adapting the technology to an experimental representative HPC system. It is understood that this assumes a degree of analysis of the existing government HPC systems of interest to identify which electrical interconnects or what category of interconnects (e.g., ones with the longest reach) are most likely to begin limiting performance as high bandwidth demands on the system continue to grow. With such information in hand a representative and scaled down HPC breadboard system is assumed as a second experimental phase that would be available for evaluating the effectiveness of the specific optical link solution.

Table 13-1 shows the recommended roadmap of specific steps to be implemented in (approximately) the next five years.

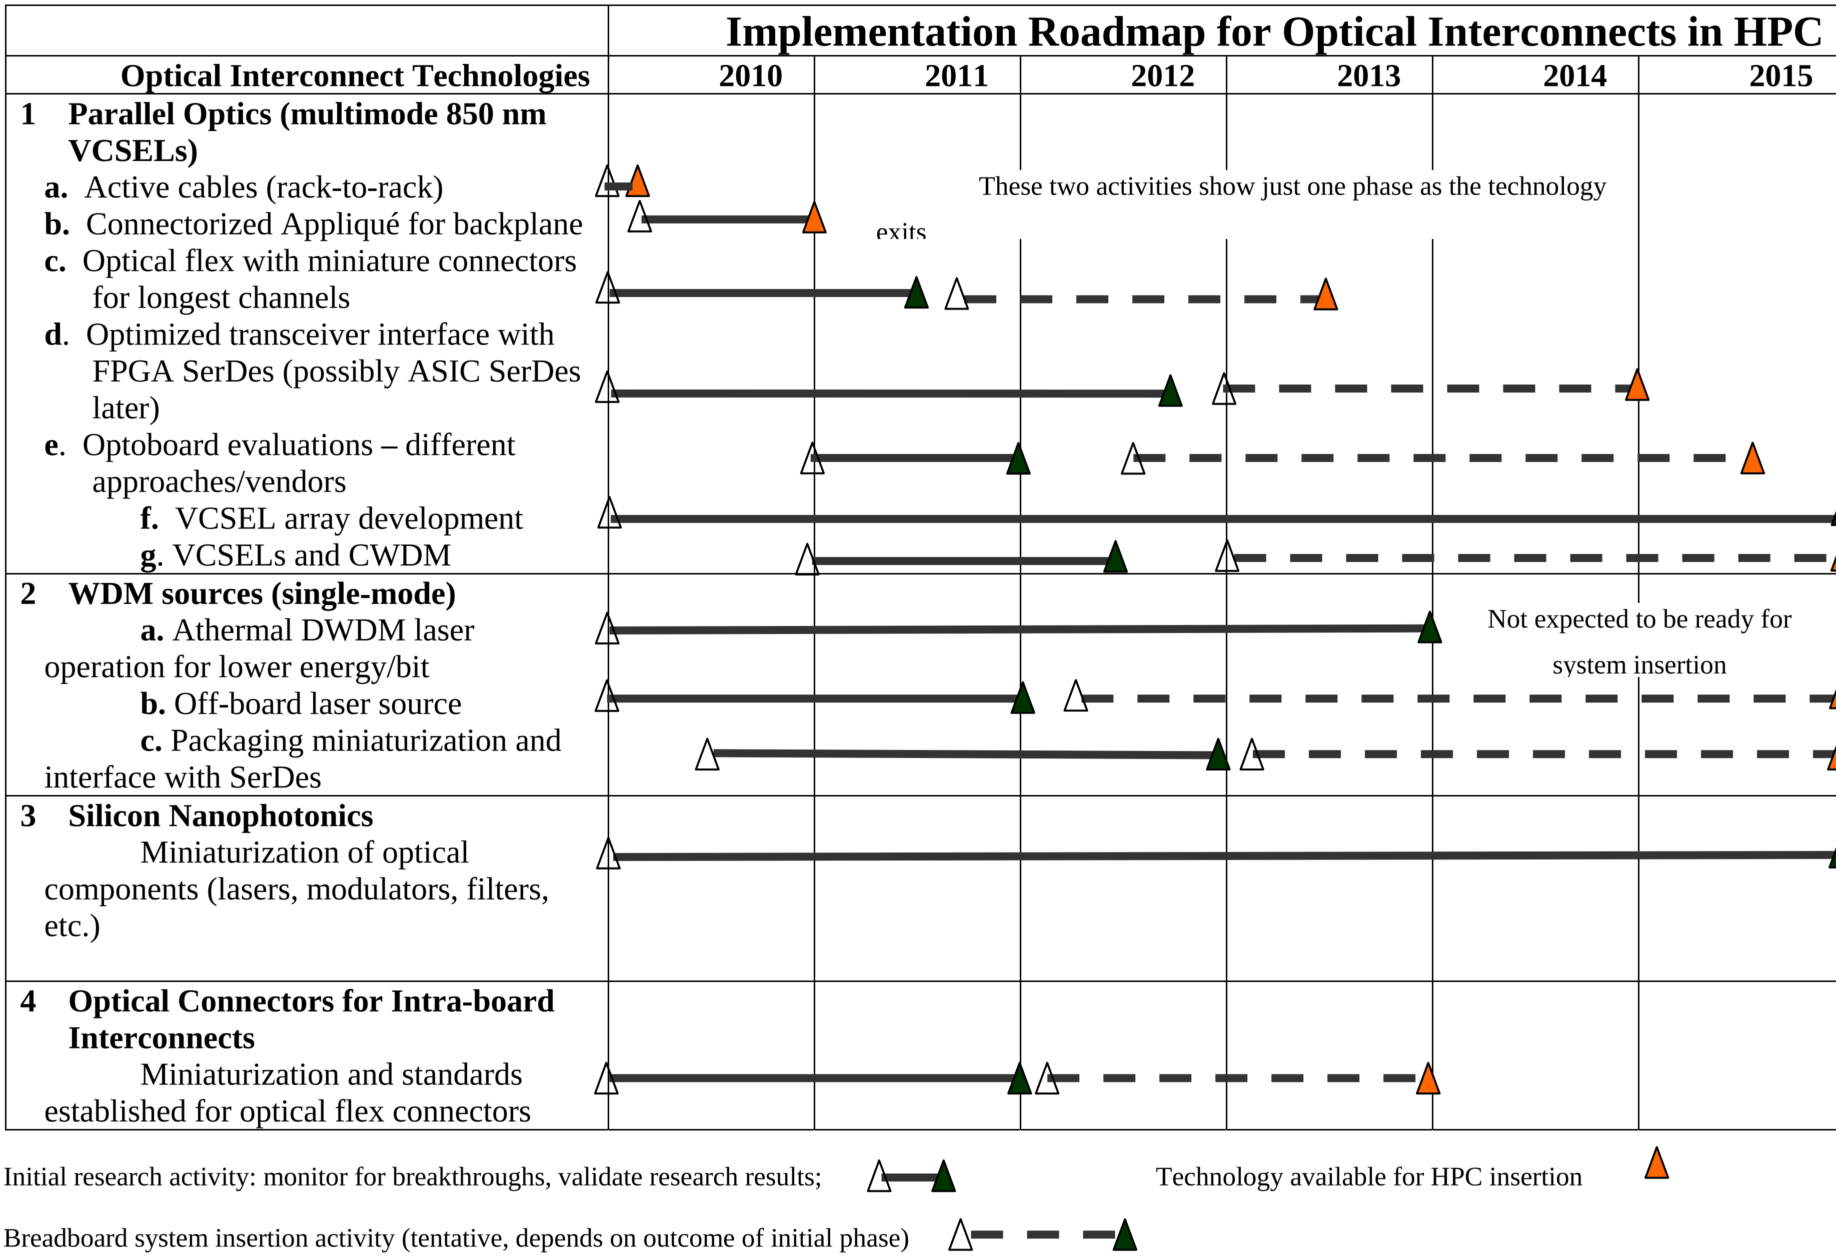

Table 13-1: Suggested Plan for Introducing Optical Interconnects into Government HPC Systems.

13.2 Implementation Steps

Brief descriptions of the list of elements shown in Table 13-1 are presented in this subsection.

13.2.1 Parallel Optics

As explained elsewhere in this Reference Guide, parallel optics refers to spatial multiplexing of multiple high speed digital communication channels implemented in the form of densely packed parallel optical waveguides and transceivers. Typically, the technology is based on multimode polymer waveguides and VCSEL light sources. It is potentially an attractive approach for interfacing to silicon ICs incorporating large numbers of SerDes I/O ports.

13.2.1.a Active Optical Cables: Active cables incorporating as many as 12 bi-directional channels each capable of supporting 10 Gbps data streams are commercially available today from several vendors (see Section 6, for a representative list of vendors). These could be inserted into an experimental HPC test bed for high speed board-to-board or rack-to-rack data transfer.

13.2.1.b Flexible Appliqué with Embedded Optical Fiber for Backplane Applications: The appliqué incorporates a network of optical fibers (with optical connector terminations) in a flexible sheet (e.g., Kapton) that can be overlaid onto a PCB or backplane). Several vendors including Molex [13-1] supply this type of optical distribution network (see Section 6 for photo). Within the bounds of vendor-supplied design rules (e.g., minimum bending radii allowed for the fibers), the user is free to specify custom routing of multiple fibers (including cross-overs). A specific optical backplane distribution network and implementation in an experimental HPC system could be realized by the end of CY 2010.

13.2.1.c Optical Flex with Miniature Connectors for Longest Channels: Several research laboratories as well as component developers are working on miniature optical waveguides that are embedded in a flexible sheet, similar to the appliqué construct, except that miniature polymer optical guides are used instead of conventional silica fiber. This offers the potential for highly dense arrays of optical waveguides that could be placed in close proximity to the high speed I/O ports of a silicon processor IC. Combined with miniature connectors (See Figure 6-1 in Section 6, for example) and small form-factor O/E transceivers, this technique could be applied to the longest intra-board links where the reach is beyond the electrical bandwidth-distance product limit. Assessing the viability of this approach in terms of determining commercial component availability (miniature transceivers, flex optical cable, miniature connectors) and component electrical and optical characterization could get started immediately and take up to 18 months for completion. Assuming a positive outcome from the first phase, implementing and testing this approach in an experimental HPC system could take another year and a half to two years.

13.2.1.d Optimized Interface of Transceiver With FPGA SerDes (Possibly ASIC SerDes Later): One of the ways to improve the efficiency of optical links is to optimize the interface between the SerDes I/O and the optical transceivers. Initially this could be demonstrated with commercial FPGAs and miniaturized E/O transceivers (an example of which is Reflex Photonics' *LightABLE* transceiver product – see Section 7, Emerging Technologies). The FPGA approach allows one to set up a series of experiments to assess the extra penalty one pays for incorporating various equalization and error correction circuitry on overall link power consumption. With an optical transceiver and short electrical connections between it and the

SerDes I/O, the equalization and error correction circuitry can be simplified or eliminated altogether for reduced power consumption. Furthermore, it is expected that the signal swing on the SerDes transmitter circuitry can be reduced and optimized to directly drive the E/O transceiver, rather than drive a lossy electrical channel. Similar power savings are expected on the receiver side. With the FPGA experiments as baseline, the work could transition to a custom ASIC replacing the FPGA. The experimental phase could begin immediately and extend to 2.5 years with an optimized implementation ready for system insertion by 2015.

13.2.1.e Optoboard Evaluations – Different Approaches and Vendors: In recent years several research efforts have investigated the possibility of incorporating optical waveguides into (more or less) conventional circuit boards such that the resulting “optoboard” would be capable of supporting not only electrical communication links but optical links as well. At least two approaches have been reported: IBM’s optoboard on FR4 [13-2] and PPC Electronic AG’s optoboard [13-3] fabricated on Schott’s Thin glass AF 32®. Both techniques embed miniature polymer optical waveguide cores into another polymer cladding layer and the combination is then laminated onto the PCB host material. In the case of PPC the glass technology is partially borrowed from the large volume LCD display market. It is not clear whether these approaches will survive beyond the research stage, however, as neither is today available as a commercial product. For this reason a one- to one and a half-year experimental phase to determine the state of development for these two approaches is recommended in the 2011 timeframe. If the technology becomes available commercially, then a three-year effort would follow to implement the optoboards in an experimental HPC system.

13.2.1.f VCSEL Array Development: In conjunction with monitoring the development status of optoboards, a parallel effort is suggested that follows the developments in arrayed VCSELs. Most data communication VCSELs today operate at 850 nm and can be directly modulated up to about 10 Gbps, although experimental devices have been modulated at much higher bit rates. Linear arrays of VCSELs are also commercially available and are typically used in active cable products. However, VCSEL research is continuing to enhance performance including efficiency, modulation rate, and low cost manufacture of linear as well as 2-D VCSEL arrays [13-4]. As shown in Table 13-1, the recommended activity is to continue monitoring VCSEL research for any breakthroughs that might occur, and where practical, evaluate promising state-of-the-art experimental devices. Although arrayed photodetectors are just as vital for an optical link at the receiver end, it seems that VCSEL technology still remains the more critical of the two and one of the main pacing items for advances in parallel optics.

13.2.1.g VCSELs and CWDM: In coarse wave division multiplexing (CWDM) the wavelengths are typically separated by 10 or 20 nm (ITU specification) in contrast to less than 1 nm in dense WDM (DWDM) systems. As discussed below, the high wavelength stability in DWDM systems requires extra power consumption to drive the temperature control electronics and the thermoelectric cooler device that maintains very precise and nearly constant laser temperature. Because of much greater wavelength spacing in CWDM systems, the TEC in most cases is unnecessary and can be eliminated, thus saving overall power. Combined with parallel multimode (MM) polymer waveguides and arrayed VCSELs, a potentially attractive, high efficiency approach for high aggregate bandwidth optical communication between HPC processor nodes presents itself. Because of MM operation alignment of the fiber and VCSEL is less critical and packaging costs can be reduced. Furthermore, it is possible to have VCSELs tuned to different wavelengths so as to achieve the benefits not only of the parallel optics, but also of wave division multiplexing (although VCSELs with wavelengths other than the standard 850 nm are much less commercially available than the tunable DFB edge emitting lasers) [13-4],[13-5]. With a broadcast element such as a star coupler, a "broadcast and select" system can also be developed. This is potentially an important area of research to consider.

In Table 13-1, an activity of one to two year's duration starting in 2011 evaluating experimental CWDM systems is indicated. The objective would be to evaluate the progress and availability of CWDM systems potentially applicable to HPC. Hardware demonstration vehicles, such as intra-board communication links, would be developed during this period. Should the technology prove sufficiently viable, a following phase would be initiated in the 2013 time frame whose objective would be to implement CWDM in a scaled, but representative version of an HPC system.

13.2.2 WDM Laser Sources and Single-Mode Waveguides

Another approach to multiplexing data is based on wavelength division multiplexing (WDM). Using edge emitters and single-mode waveguides and fiber, WDM is extensively used in the telecommunications industry because of several unique characteristics that could likewise prove to be advantageous in HPC systems. Among these is the ability to have simultaneous channels multiplexed onto the same fiber. For dense WDM (DWDM) today's ITU standards dictate a number of allowed wavelength spacings including 200 GHz, 100 GHz, 50 GHz, and possibly 25 GHz and 12.5 GHz for future systems. For example, 100 GHz spacing will accommodate 40 wavelengths in C-band (1530 to 1565 nm), all on one fiber. This is very beneficial in a dense interconnect environment for drastically reducing the number of physical connectors. Another unique feature of single-mode systems is the ability to broadcast a data stream to many destinations, for example, by means of a broadcast element such as a star coupler. However, there are many technical issues that need to be resolved before this technology approach can be successfully applied to HPC systems including packaging miniaturization and improved overall efficiency. The following chronological list of suggested activities follows the corresponding Implementation Roadmap items shown in Table 13-1 and is intended to be a guide for introducing this facet of optical technology into HPC systems.

13.2.2.a Athermal DWDM Laser Operation for Lower Energy/Bit: As discussed elsewhere in this Reference Guide, DWDM requires that the edge-emitting lasers (usually distributed feedback (DFB)) be ultra stable with respect to wavelength. This implies use of thermoelectric coolers (TECs) and control electronics to maintain the laser temperature to within several tenths of a degree Centigrade. The extra power needed for the TEC can be a substantial fraction of the total laser module power consumption (60% to over 90% for a DFB laser transmitter assembly [13-6]). In telecommunications laser transmitters, the electrical power to drive a 20 mW DFB laser is 300 mW while the power required for the TEC cooler can be as large as 10W [13-7]! However, techniques are being explored in research laboratories to do away with the TEC to reduce power consumption and yet maintain a viable DWDM system. Some of these techniques involve software that tracks the wavelength as the ambient temperature changes and adjusts local heaters on the laser chip to compensate for the temperature change [13-7]. These local on-chip heaters require significantly less power than a much larger TEC device that is incorporated into the laser transmitter assembly. Although currently these are strictly exploratory techniques, it is important to follow this technology, should some breakthrough occur in significantly lowering the energy-per-bit in optical links using DWDM. As shown in the Table, this activity should be followed for the next several years.

13.2.2.b Off-Board DFB Laser Source: In contrast to direct modulation, it is possible to use a CW laser and modulate indirectly using an external modulator. In fact, many InP-based laser transmitters for telecommunications employ an electro-absorption modulator that is monolithically integrated with the laser that operates in CW mode. However, it is also possible to have the CW laser located remotely, or off-board, and to route the light to several independent modulators using an optical power splitter. Each modulator works independently of the others and is driven with its own data stream. Other than the remote laser location, this scheme is used in Luxtera's active cable product where the DFB laser chip is mounted directly onto the CMOS chip incorporating the silica waveguides and 1:4 power splitter as well as the drive electronics (Blazar [13-8]). Several lasers with different wavelengths can service several sets of modulators. Essentially, this scheme is a combination of spatial and wavelength division multiplexing. Keeping the laser off-board has several advantages: a) laser temperature control can be localized and made independent of the temperature at the node processor site; b) CW laser operation can be optimized for power, low relative intensity noise (RIN), and high efficiency without regard to high modulation speeds; c) Maintainability is easier since a suitable location can be selected to serve this purpose. A disadvantage is in case of failure multiple data streams are affected. One would need to have backup, redundant laser sources and automatic switching to a working laser in case of failure.

In Table 13-1 a two-year task is recommended to monitor the research progress in this area starting in early 2010. Should this approach prove to be viable for HPC, a four year effort is suggested for implementation in an experimental HPC test bed.

13.2.2.c Packaging miniaturization and interface with SerDes : Just as in the VCSEL case (subsection 13.2.1.d), if DFB edge-emitting lasers are to be viable for HPC application, they need to be intimately packaged with the electronics. Specifically, they need to be packaged with the high speed I/O SerDes circuitry (except for the case where they are purposely located off-board and used with external modulators as described in the previous subsection). This will most likely involve vertical stacking of electronic and photonic chips in order to achieve the highest speed performance and signal integrity of the optically modulated signals. On the receiver side the same is true with the photodetectors, i.e., they too must be connected with minimum circuit parasitics to their respective TIAs and associated electronics to achieve the highest performance.

This important research area should be monitored for the next two years or more to compare its progress relative to other approaches shown in the Table. Should significant headway be made in this area implementation into an experimental test bed should begin no later than early 2013 for the technology to be ready for HPC insertion by 2016.

13.2.2.d Silicon Nanophotonics: As the name implies this area of research aims at miniaturizing the optical components that can be fabricated monolithically on a silicon substrate to where the functionality and optical component density is on par with that of the electrical integrated circuit. Some examples are given in Section 7, Emerging Technology. Successful realization of this far-reaching technology would allow data transfer in the optical domain to dominate computer interconnects. Inherently linked to this technology is reduction of overall power consumption, since component miniaturization implies smaller areas for heat transfer and without commensurate power reduction power densities and device temperatures would increase to unacceptable levels. A key requirement is a silicon-compatible laser source. Micron-sized semiconductor lasers are under development in many university and industrial research laboratories [13-9], [13-10]. However, the transition from research studies to actual practical implementation is seen to be at least seven years away [13-10]. For this reason the corresponding activity shown in Table 13-1 is only indicated as taking place in the first phase, wherein the research is monitored for breakthroughs, and if practical, experimental devices are evaluated. The technology is not expected to transition to any insertion into an HPC breadboard system in the next five years.

13.2.2.e Optical Connectors: Although there is a wealth and diversity of optical connectors for long- and short-haul telecommunication applications as well as local rack-to-rack optical interconnects for small and large data centers (see Section 4), there is much less available choice when it comes to miniature optical connectors for intra-board interconnects. It is anticipated that before optoboards become available (assuming they become commercially viable), an interim solution will be to employ short intra-board optical flex circuits that are compatible with interconnecting on-board processing nodes. As described in Sections 7, flex circuits containing multiple parallel fibers or optical polymer waveguides can be used to interconnect intra-board nodes for those cases where the reach is beyond the limit for an electrical connection. Compatible optical connectors are required. Experimental versions of such connectors have been described at several research laboratories [13-11] (see Section 6, Existing Technologies), but to the best of our knowledge, commercial availability and standards are still lacking for such products.

As the last row in Table 13-1 indicates, an experimental phase to sample research connectors and to monitor progress in this technology area is recommended for about two years starting immediately. A follow-on phase to implant the connectors with flex optical circuits in an experimental HPC test bed is suggested for a one and half-year effort to be ready for system insertion by 2014.

13.3 References

[13-1] Molex Website: http://www.molex.com/molex/products/datasheet.jsp?part=active/1064010000_FIBER_OPTIC_ASSEMB.xml&channel=Products&Lang=en-US

[13-2] Doany, F., et al.: 160 Gb/s Bidirectional Polymer-Waveguide Board-Level Optical Interconnects Using CMOS-Based Transceivers. IEEE Transactions on Advanced Packing, 32(2):345-359 (May) 2009.

[13-3] PPC Electronic AG Website: http://www.ppc-electronic.com/

[13-4] Hoffmann, W., et al.: One and Two-Dimensional Long-Wavelength VCSEL Arrays for DWDM Applications and Optical Interconnects. IEEE, 21st International Semiconductor Laser Conference, ISLC, paper ThA4, September 2008.

[13-5] Patel, R. R., et al.: Multiwavelength Parallel Optical Interconnects for Massively Parallel Processing. IEEE Journal of Selected Topics in Quantum Electronics, 9:657-666 (March/April) 2003.

[13-6] Private Communication: Dr. Gregory Whaley, Lockheed Martin, Eagan, MN.

[13-7] Infinera Technical Memorandum: Uncooled, Full C-band, Multi-Wavelength DWDM Transmitter enabled by Adaptive Self Monitoring: Phase II. DARPA SBIR, W31P4Q-05-C-0302 Scientific and Technical Final Report, July 31, 2008

[13-8] Luxtera Article: Blazar 40 Gbps Optical Active Cable. Luxtera Website: http://www.luxtera.com/white-papers.html

[13-9] Cornell-Norfolk State-Purdue University Research Report Described At: http://www.powermanagementdesignline.com/news/219501015

[13-10] Private Communication: Dr. Raymond Beausoleil, Hewlett-Packard Laboratories (Palo Alto).

[13-11] Private Communication: Dr. Bert Offrein, IBM Zurich Research Laboratories.

14 Glossary

ADP	Ammonia Dihydrogen Phosphate ($NH_4H_2PO_4$) – A material used to make Pockels cell modulators.
AMI [14-1]	Alternate Mark Inversion – When used on a T-carrier, the code is known as Alternate Mark Inversion because, in this context, a binary 1 is referred to as a "mark", while a binary 0 is called a "space". The coding was used extensively in first-generation PCM networks, and is still commonly seen on older multiplexing equipment today, but successful transmission relies on no long runs of zeroes being present. For voice traffic, if the encoded value of the voice sample is zero (00000000), the code 00000010 is sent instead. This ensures that no more than 15 consecutive zeros will ever be sent, thus ensuring synchronization. The modification of bit 7 causes a change to voice that is undetectable by the human ear, but it is an unacceptable corruption of a data stream. Data channels are required to use some other form of pulse-stuffing, such as always setting bit 8 to 1, in order to maintain one's density. Of course, this lowers the effective data throughput to 56 kbit/s per channel.
AP	Alternating-Phase
APD [14-1]	Avalanche Photodiode – APDs are photodetectors that can be regarded as the semiconductor analog to photomultipliers. By applying a high reverse bias voltage (typically 100-200 V in silicon), APDs show an internal current gain effect (around 100) due to impact ionization (avalanche effect). However, some silicon APDs employ alternative doping and beveling techniques compared to traditional APDs that allow greater voltage to be applied (> 1500 V) before breakdown is reached and hence a greater operating gain (> 1000).
AWG [14-1]	Array Waveguide Grating – AWGs are commonly used as optical (de)multiplexers in wavelength division multiplexed (WDM) systems. These devices are capable of multiplexing a large number of wavelengths into a single optical fiber, thereby increasing the transmission capacity of optical networks considerably.
BER [14-1]	Bit Error Rate – In telecommunication, an error ratio is the ratio of the number of bits, elements, characters, or blocks incorrectly received to the total number of bits, elements, characters, or blocks sent during a specified time interval. The most commonly encountered ratio is the bit error ratio (BER) - also sometimes referred to as bit error rate.

BPSK [14-1]	Binary Phase Shift Keying – BPSK is the simplest form of PSK. It uses two phases which are separated by 180° and so can also be termed 2-PSK. It does not particularly matter exactly where the constellation points are positioned, and in this figure they are shown on the real axis, at 0° and 180°. This modulation is the most robust of all the PSKs since it takes serious distortion to make the demodulator reach an incorrect decision. It is, however, only able to modulate at 1 bit/symbol (as seen in the figure) and so is unsuitable for high data-rate applications when bandwidth is limited.
BW	Bandwidth
CRZ	Chirped Return-to-Zero
CSRZ [14-1]	Carrier-Suppressed Return-to-Zero – CSRZ is an optical signal format. In CSRZ the field intensity drops to zero between consecutive bits (RZ), and the field phase alternates by π between neighboring bits, so that if the phase of the signal is e.g. 0 in even bits (bit number 2**n**), the phase in odd bit slots (bit number 2**n**+1) will be π, the phase alternation amplitude. In its standard form CSRZ is generated by a single Mach-Zehnder Modulator (MZM), driven by two sinusoidal waves at half the bit rate B_R, and in phase opposition. This gives rise to characteristically broad pulses (duty cycle 67%).
CW [14-1]	Continuous Wave – A continuous wave or continuous waveform is an electromagnetic wave of constant amplitude and frequency; and in mathematical analysis, of infinite duration. Continuous wave is also the name given to an early method of radio transmission, in which a carrier wave is switched on and off. Information is carried in the varying duration of the on and off periods of the signal. In radio transmission, CW waves are also known as "undamped waves", to distinguish this method from damped wave transmission.
CWDM [14-1]	Coarse Wave Division Multiplexing – Originally, CWDM was fairly generic and meant a number of different things. In general, these things shared the fact that the choice of channel spacings and frequency stability was such that Erbium Doped Fiber Amplifiers (EDFAs) could not be utilized. Prior to the relatively recent ITU standardization of the term, one common meaning for Coarse WDM meant two (or possibly more) signals multiplexed onto a single fiber, where one signal was in the 1550-nm band, and the other in the 1310-nm band. Recently the ITU has standardized a 20 nm channel spacing grid for use with CWDM, using the wavelengths between 1310 nm and 1610 nm. Many CWDM wavelengths below 1470 nm are considered "unusable" on older G.652 specification fibers, due to the increased attenuation in the 1310-1470 nm bands. Newer fibers which conform to the G.652.C and G.652.D standards, such as Corning SMF-28e and Samsung Widepass nearly eliminate the "water peak" attenuation peak and allow for full operation of all twenty ITU CWDM channels in metropolitan networks.

DBPSK [14-1]	Differential BPSK – For BPSK and QPSK there is an ambiguity of phase if the constellation is rotated by some effect in the communications channel the signal passes through. This problem can be overcome by using the data to **change** rather than **set** the phase. For example, in differentially-encoded BPSK a binary '1' may be transmitted by adding 180° to the current phase and a binary '0' by adding 0° to the current phase. This kind of encoding may be demodulated in the same way as for non-differential PSK but the phase ambiguities can be ignored. Thus, each received symbol is demodulated to one of the M points in the constellation and a comparator then computes the difference in phase between this received signal and the preceding one. The difference encodes the data as described above.
DBR [14-1]	Distributed Bragg Reflector – A DBR is a high quality reflector used in waveguides, such as optical fibers. It is a structure formed from multiple layers of alternating materials with varying refractive index, or by periodic variation of some characteristic (such as height) of a dielectric waveguide, resulting in periodic variation in the effective refractive index in the guide. Each layer boundary causes a partial reflection of an optical wave. For waves whose wavelength is close to four times the optical thickness of the layers, the many reflections combine with constructive interference, and the layers act as a high-quality reflector. The range of wavelengths that are reflected is called the photonic stopband. Within this range of wavelengths, light is "forbidden" to propagate in the structure.
DFB [14-1]	Distributed Feedback Laser – A DFB is a type of laser diode where the active region of the device is structured as a diffraction grating. The grating, known as a distributed Bragg reflector, provides optical feedback for the laser due to Bragg scattering from the structure. Since the grating provides feedback, DFB lasers do not use discrete mirrors to form the optical cavity (as are used in conventional laser designs). The grating is constructed so as to reflect only a narrow band of wavelengths, and thus produce a narrow linewidth of laser output. Altering the temperature of the device causes the pitch of the grating to change due to thermal expansion. This alters the reflection wavelength of the grating structure and thus the wavelength of the laser output, producing a tunable laser. The tuning range is usually of the order of 6 nm for a ~50 K change in temperature. Altering of the modulation rate of the current powering the laser will also tune the device. DFB lasers are often used in optical communication applications such as DWDM where a tunable laser signal is desired.
DML [14-2]	Directly-Modulated Laser – The light output is controlled by a current applied directly to the device.

DOS [14-3]	Digital Optic Switch – A Y-coupler modified by the addition of electrodes which can impose an electric field on the waveguide material. The material can be lithium niobate which changes its refractive index under the influence of an electric field.
DQPSK [14-1]	Differential QPSK – For BPSK and QPSK there is an ambiguity of phase if the constellation is rotated by some effect in the communications channel the signal passes through. This problem can be overcome by using the data to **change** rather than **set** the phase. For example, in differentially-encoded QPSK, the phase-shifts are 0°, 90°, 180°, -90° corresponding to data '00', '01', '11', '10'. This kind of encoding may be demodulated in the same way as for non-differential PSK but the phase ambiguities can be ignored. Thus, each received symbol is demodulated to one of the M points in the constellation and a comparator then computes the difference in phase between this received signal and the preceding one. The difference encodes the data as described above.
DSF [14-1]	Dispersion-Shifted Fiber – DSF, specified in ITU-T G.653, is a type of single-mode optical fiber with a core-clad index profile tailored to shift the zero-dispersion wavelength from the natural 1300 nm in silica-glass fibers to the minimum-loss window at 1550 nm. The group-velocity or **intramodal** dispersion which dominates in single-mode fibers is comprised of both material and waveguide dispersion. Waveguide dispersion can be made more negative by changing the index profile and thus be used to offset the fixed material dispersion, shifting or flattening the overall intramodal dispersion. This is advantageous because it allows a communication system to possess both low dispersion and low attenuation. However, when used in wavelength division multiplexing systems, dispersion-shifted fibers can suffer from four-wave mixing which causes intermodulation of the independent signals. As a result nonzero dispersion shifted fiber is often used.

DWDM [14-1]	Dense Wave Division Multiplexing – DWDM refers originally to optical signals multiplexed within the 1550-nm band so as to leverage the capabilities (and cost) of erbium doped fiber amplifiers (EDFAs), which are effective for wavelengths between approximately 1525 nm - 1565 nm (C band), or 1570 nm - 1610 nm (L band). EDFAs were originally developed to replace SONET/SDH optical-electrical-optical (OEO) regenerators, which they have made practically obsolete. EDFAs can amplify any optical signal in their operating range, regardless of the modulated bit rate. In terms of multi-wavelength signals, so long as the EDFA has enough pump energy available to it, it can amplify as many optical signals as can be multiplexed into its amplification band (though signal densities are limited by choice of modulation format). EDFAs therefore allow a single-channel optical link to be upgraded in bit rate by replacing only equipment at the ends of the link, while retaining the existing EDFA or series of EDFAs along a long haul route. Furthermore, single-wavelength links using EDFAs can similarly be upgraded to WDM links at reasonable cost. The EDFAs cost is thus leveraged across as many channels as can be multiplexed into the 1550-nm band.
EAM [14-1]	Electro-Absorption Modulator – A semiconductor diode that modulates light from a laser that is separate from it, but may be fabricated on the same wafer. Turning the current on causes absorption of the light.
EDFA [14-1]	Erbium Doped Fiber Amplifier – The EDFA is the most deployed fiber amplifier as its amplification window coincides with the third transmission window of silica-based optical fiber. Two bands have developed - C-band from approximately 1525 nm - 1565 nm, and L-band from approximately 1570 nm to 1610 nm. Both of these bands can be amplified by EDFAs, but it is normal to use two different amplifiers, each optimized for one of the bands. EDFAs have two commonly-used pumping bands - 980 nm and 1480 nm.
EDWA [14-1]	Erbium Doped Waveguide Amplifier – An EDWA is an optical amplifier that uses a waveguide to boost an optical signal, analogous to an EDFA.
EEL	Edge-Emitting Laser – A semiconductor diode laser whose light is emitted from the edge of the laser.
E_g [14-1]	Bandgap Energy – In solid state physics and related applied fields, the band gap, also called an energy gap or stop band, is a region where a particle or quasiparticle is forbidden from propagating. For insulators and semiconductors, the band gap generally refers to the energy difference between the top of the valence band and the bottom of the conduction band.

EIA [14-1]	Electronic Industries Alliance – The EIA (until 1997 it was the Electronic Industries Association) is a trade organization composed as an alliance of trade associations for electronics manufacturers in the United States. Those associations in turn govern sectors of EIA standards activity. The associations are: • Formerly CEA – The Consumer Electronics Association (As of January 1, 2005 CEA withdrew from its sector affiliation with EIA) • ECA – The Electronic Components, Assemblies, and Materials Association • GEIA – The Government Electronics and Information Technology Association • JEDEC – The JEDEC Solid State Technology division, formerly Joint Electron Devices Engineering Councils • TIA – The Telecommunications Industry Association EIA is accredited by ANSI to help develop standards on electronic components, consumer electronics, electronic information, telecommunications, and Internet security. The recommended standards (currently EIA-#) are designed so that manufacturer's equipment can be interchanged and compatible.
EML	Electro-Absorption Modulated Laser – A laser, typically a DFB, with an integrated electro-absorption modulator.
E-O	Electro-Optic
ER [14-1]	Extinction Ratio – In telecommunications, extinction ratio ($\mathbf{r}_e$) is the ratio of two optical power levels, of a digital signal generated by an optical source, **e.g.,** a laser diode, where $\mathbf{P}_1$ is the optical power level generated when the light source is "on," and $\mathbf{P}_2$ is the power level generated when the light source is "off." **Note:** The extinction ratio may be expressed as a fraction or in dB.
ESD [14-1]	Electrostatic Discharge – ESD is the sudden and momentary electric current that flows between two objects at different electrical potentials. The term is usually used in the electronics and other industries to describe momentary unwanted currents that may cause damage to electronic equipment.
ETDM	Electrical Time Division Multiplexing
FC [14-1]	Fibre Channel – Fibre Channel is a gigabit-speed network technology primarily used for storage networking. Fibre Channel is standardized in the T11 Technical Committee of the InterNational Committee for Information Technology Standards (INCITS), an American National Standards Institute (ANSI)–accredited standards committee. It started for use primarily in the supercomputer field, but has become the standard connection type for storage area networks (SAN). Fibre Channel signaling can run on both twisted pair copper wire and fiber-optic cables.

FOCIS [14-1]	Fiber Optic Connector Intermateability Standard – FOCIS specifications for cable connectors from the TIA that define the requirements for interconnection between fiber-optic plugs and sockets.
FOTP	Fiber Optic Test Procedures
FP [14-1]	Fabry-Perot – An optical structure containing a pair of mirrors at opposite ends of a cavity. Light reflects back and forth between the mirrors, and one or both transmit a fraction of the resonant frequency. The resonance is created by making the distance of one round trip between mirrors equal to an integral number of wavelengths of the cavity material. Optically speaking, this is an interferometer, because it relies on the interference of light for its operation. A Fabry-Perot device becomes a light generator when a laser medium is used in the cavity; otherwise it is a passive filter.
FPGA	Field Programmable Gate Array
FSK [14-1]	Frequency Shift Keying – A simple digital modulation technique that uses two frequencies for 0 and 1.
FWHM [14-1]	Full Width at Half Maximum – FWHM is an expression of the extent of a function, given by the difference between the two extreme values of the independent variable at which the dependent variable is equal to half of its maximum value.
FWM [14-1]	Four-Wave Mixing – Four-wave mixing is an intermodulation distortion in optical systems, similar to the third order intercept in electrical systems.
Ge [14-1]	Germanium – This is a lustrous, hard, silver-white metalloid that is chemically similar to tin. Ge forms a large number of organometallic compounds and is an important semiconductor material used in transistors and fiber optic communication networks. Unlike most semiconductors, germanium has a small band gap, allowing it to efficiently respond to infrared light.
HMG	Hollow Metal Waveguide

HPC [14-1]	High Performance Compute – HPC refers to the use of (parallel) supercomputers and computer clusters, that is, computing systems comprised of multiple (usually mass-produced) processors linked together in a single system with commercially available interconnects. This is in contrast to mainframe computers, which are generally monolithic in nature. While a high level of technical skill is undeniably needed to assemble and use such systems, they can be created from off-the-shelf components. Because of their flexibility, power, and relatively low cost, HPC systems increasingly dominate the world of supercomputing. Usually, computer systems in or above the teraflop-region are counted as HPC-computers. The term is most commonly associated with computing used for scientific research. A related term, High-performance technical computing (HPTC), generally refers to the engineering applications of cluster-based computing (such as computational fluid dynamics and the building and testing of virtual prototypes). Recently, HPC has come to be applied to business uses of cluster-based supercomputers, such as data warehouses, line-of-business (LOB) applications and transaction processing.
I	In-Phase – The in-phase component of a quadrature phase shift keyed signal and also the in-phase axis on a constellation diagram.
IC [14-1]	Integrated Circuit – In electronics, an IC (also known as microcircuit, microchip, silicon chip, or chip) is a miniaturized electronic circuit (consisting of semiconductor devices, as well as passive components) that has been manufactured in the surface of a thin substrate of semiconductor material.
ISI [14-1]	Intersymbol Interference – In telecommunication, ISI means a form of distortion of a signal that causes the previously transmitted symbols to have an effect on the currently received symbol. This is usually an unwanted phenomenon as the previous symbols have similar effect as noise, thus making the communication less reliable. ISI is usually caused by echoes or non-linear frequency response of the channel. Ways to fight against intersymbol interference include adaptive equalization or error correcting codes.
ITU [14-1]	International Telecommunications Union – The ITU (French: **Union internationale des télécommunications**,) is an international organization established to standardize and regulate international radio and telecommunications. It was founded as the **International Telegraph Union** in Paris on May 17, 1865. Its main tasks include standardization, allocation of the radio spectrum, and organizing interconnection arrangements between different countries to allow international phone calls — in which regard it performs for telecommunications a similar function to what the UPU performs for postal services. It is one of the specialized agencies of the United Nations, and has its headquarters in Geneva, Switzerland, next to the main United Nations campus.

KDP [14-1]	Potassium Dihydrogen Phosphate (KH_2PO_4) – (also monopotassium phosphate, or monobasic potassium phosphate, MKP) As a crystal, it is noted for its non-linear optical properties. Used in optical modulators and for non-linear optics such as SHG (second harmonic generation).
LED [14-1]	Light Emitting Diode – An LED is a semiconductor diode that emits incoherent narrow-spectrum light when electrically biased in the forward direction of the p-n junction. This effect is a form of electroluminescence. An LED is usually a small area source, often with extra optics added to the chip that shapes its radiation pattern. The color of the emitted light depends on the composition and condition of the semiconducting material used, and can be infrared, visible, or near-ultraviolet.
$LiNiO_3$ [14-1]	Lithium Niobate ($LiNbO_3$) – Lithium niobate is a compound of niobium, lithium, and oxygen. It is a colorless solid that is insoluble in water. Its melting point is 1257 °C and its density is 4.65 g/cm³. It is transparent for wavelengths between 350 and 5200 nanometers, and has a bandgap of around 4 eV. Single monocrystals are used in laser frequency doubling, nonlinear optics, Pockels cells, optical parametric oscillators, Q-switching devices for lasers, other acousto-optic devices, optical switches for gigahertz frequencies, etc. It is an excellent material for manufacture of optical waveguides. Lithium niobate is used extensively in the telecoms market, eg. in the mobile telephones and optical modulators. It is the material of choice for the manufacture of surface acoustic wave devices.
LSB	Least Significant Bit – Bit 0 in an n bit binary number, the bit with the least weight (2^0). The last or rightmost bit when the number is written in the usual way.
MAN [14-1]	Metro Area Network – A data network intended to serve an area the size of a large city. Such networks are being implemented by innovative techniques, such as running optical fiber through subway tunnels.
MEMS [14-1]	Microelectromechanical Systems – MEMS is the technology of the very small, and merges at the nano-scale into nanoelectromechanical systems (NEMS) and Nanotechnology. MEMS are also referred to as micro machines, or **Micro Systems Technology** (**MST**). MEMS are separate and distinct from the hypothetical vision of Molecular nanotechnology or Molecular Electronics. MEMS generally range in size from a micrometer (a millionth of a meter) to a millimeter (thousandth of a meter). At these size scales, the standard constructs of classical physics do not always hold true. Due to MEMS' large surface area to volume ratio, surface effects such as electrostatics and wetting dominate volume effects such as inertia or thermal mass. Finite element analysis is an important part of MEMS design.
MM	Multimode

MMF [14-1]	Multimode Fiber – An optical fiber with a larger core than single-mode fiber. It is the most commonly used fiber for short distances such as LANs. Light can enter the core at different angles, making it easier to connect the light source to broader light sources such as LEDs. However, light rays travel down multiple reflective paths (modes) in multimode fiber, causing modal dispersion, which is a broadening of the pulses at the receiving end.
MSA	Multi-Source Agreements – Agreement among vendors to adhere to a standard that they have developed to promote commercialization of a product.
MSB [14-1]	Most Significant Bit – Bit n-1 in an n bit binary number, the bit with the greatest weight (2^(n-1)). The first or leftmost bit when the number is written in the usual way.
MSM	Metal-Semiconductor-Metal
MZ	Mach-Zehnder
MZI [14-1]	MZ Interferometer – The MZI (named after physicists Ludwig Mach and Ludwig Zehnder) is a device used to determine the phase shift caused by a small sample which is placed in the path of one of two collimated beams (thus having plane wavefronts) from a coherent light source.
MZM [14-1]	MZ Modulator – An intensity modulator based on an MZI converts changes in phase (which we can create directly) to changes in signal amplitude. The basic principle is that you have a balanced configuration of a splitter and a combiner connected by a pair of matched waveguides. When something is done to create a phase difference between the signals in the two matched waveguides interference in the recombination process causes differences in the amplitude of the output signal.
NA [14-1]	Numeric Aperture – In optics, the NA of an optical system is a dimensionless number that characterizes the range of angles over which the system can accept or emit light. The exact definition of the term varies slightly between different areas of optics.
NDSF	Non-Dispersion-Shifted Fiber – “Regular” fiber
NEC [14-1]	National Electric Code – The NEC, or NFPA 70, is a U.S. standard for the safe installation of electrical wiring and equipment. It is part of the National Fire Codes series published by the National Fire Protection Association (NFPA). "National Electrical Code" and "NEC" are registered trademarks of the NFPA. While the NEC is not itself a U.S. law, NEC use is commonly mandated by state or local law, as well as in many jurisdictions outside of the United States. The NEC codifies the requirements for safe electrical installations into a single, standardized source.

NRZ [14-1]	Non-Return to Zero – In telecommunication, an NRZ line code is a binary code in which "1s" are represented by one significant condition and "0s" are represented by the other significant condition, with no other neutral or rest condition. The pulses have more energy than a RZ code. Unlike RZ, NRZ does not have a rest state. NRZ is not inherently a self-synchronizing code, so some additional synchronization technique (perhaps a run length limited constraint or a parallel synchronization signal) must be used to avoid bit slip.
NRZI [14-1]	Non-Return to Zero Inverted – NRZI is a method of mapping a binary signal to a physical signal for transmission over some transmission media. The two level NRZI signal has a transition at a clock boundary if the bit being transmitted is a logical one, and does not have a transition if the bit being transmitted is a logical zero. "One" is represented by a transition of the physical level. "Zero" has no transition. The transition occurs on the leading edge of the clock for the given bit. However, even NRZI can have long series of zeros (or ones if transitioning on "zero"), so clock recovery can be difficult unless some other encoding is used on top of it. NRZI encoding is used for recording on magnetic tape and for data transmission in the standard USB.
NZ-DSF [14-1]	Non-Zero Dispersion Shifted Fiber – NZ-DSF, specified in ITU-T G.655, is a type of single-mode optical fiber which was designed to overcome the problems of dispersion-shifted fiber. NZDSF is available in two primary flavors: NZD+ and NZD-, which differ in their zero-dispersion wavelengths. These are typically around 1510 nm and 1580 nm, respectively. Because the zero-dispersion point of NZDSF is outside of the normal communications window, four-wave mixing and other non-linear effects are minimized. Other types of NZDSF include RS-NZDSF which has a reduced slope in its change of dispersion and large core NZDSF which further reduces residual non-linear distortion under high launch power. Some long-haul fiber paths will alternate NZD+ and NZD- segments to provide self-dispersion compensation with uniformly low dispersion across the minimum-loss window at 1550 nm.
OADM [14-1]	Optical Add/Drop Multiplexer – An OADM is a device used in wavelength-division multiplexing systems for multiplexing and routing different channels of light into or out of an SMF. They are generally used for the construction of ring-based optical networks. "Add" and "drop" here refer to the capability of the device to add one or more new wavelength channels to an existing multi-wavelength WDM signal, and/or to drop (remove) one or more channels, routing those signals to another network path. An OADM may be considered to be a specific type of optical cross-connect.
OEIC	Optoelectronic Integrated Circuit
OFSTP	Optical Fiber Standard Test Procedures

OH [14-1]	Hydroxyl – Hydroxyl in chemistry stands for a molecule consisting of an oxygen atom and a hydrogen atom connected by a covalent bond. The neutral form is a hydroxyl radical and the hydroxyl anion is called a hydroxide.
OMA [14-1]	Optical Modulation Amplitude – In telecommunications, OMA is the difference between two optical power levels, of a digital signal generated by an optical source, **e.g.,** a laser diode, where $\mathbf{P}_1$ is the optical power level generated when the light source is "on," and $\mathbf{P}_0$ is the power level generated when the light source is "off."
OOK [14-1]	On-Off Keying – OOK is a type of modulation that represents digital data as the presence or absence of a carrier wave. In its simplest form, the presence of a carrier for a specific duration represents a binary one, while its absence for the same duration represents a binary zero. Some more sophisticated schemes vary these durations to convey additional information. OOK is not very spectrally efficient due to the abrupt changes in amplitude of the carrier wave. At low to medium signalling speeds, this can be mitigated by adjusting the rise and fall rates of the carrier's amplitude. At high speeds, more efficient modulation modes (such as frequency-shift keying) are normally used instead.
OSA	Optical Sub-Assembly
OTDM	Optical Time Division Multiplexing
OXC [14-1]	Optical Cross-Connect Switch – An OXC is device used by telecommunications carriers to switch high-speed optical signals in a fiber optic network.
PC	Physical Contact
PCB [14-1]	Printed Circuit Board – PCBs are used to mechanically support and electrically connect electronic components using conductive pathways, or **traces**, etched from copper sheets laminated onto a non-conductive **substrate**. Alternative names are printed wiring board (PWB), and etched wiring board. **Populating** the board with electronic components forms a printed circuit assembly (PCA), also known as a printed circuit board assembly (PCBA).
PD [14-1]	Photodiode – A PD is a semiconductor diode that functions as a photodetector. Photodiodes are packaged with either a window or optical fiber connection, in order to let in the light to the sensitive part of the device.

PIC [14-1]	Photonic Integrated Circuit – A PIC is a device that integrates multiple photonic functions and as such is analogous to an Electronic Integrated Circuit. However the major difference between the two being that a photonic integrated circuit provides functionality for information signals imposed on optical wavelengths typically in the Visible spectrum or near infrared 850nm-1650nm. Photonic integrated circuits can allow optical systems to be made more compact and higher performance than with discrete optical components. They also offer the possibility of integration with electronic circuits to provide increased functionality.
PLC	Planar Lightwave Circuit – The same as a PIC except with a silicon substrate and silica waveguides.
PM	Polarization Maintaining
PMF [14-1]	Polarization Maintaining Fiber – In fiber optics, PMF (or PM fiber) is an optical fiber in which the polarization planes of light waves launched into the fiber are maintained during propagation with little or no cross-coupling of optical power between the polarization modes. The maintaining of polarization is increasingly important in modern telecommunication systems, as our society demands higher bandwidth and more complex optical networks. Several different designs of PM fiber are used. Most work by inducing stress in the core via a non-circular cladding cross-section, or via rods of another material included within the cladding. Several different shapes of rod are used, and the resulting fiber is sold under brand names such as "Panda" and "Bow-tie". The differences in performance between these types of fiber are subtle. Polarization-maintaining optical fibers are used in special applications, such as in fiber optic sensing, interferometry and quantum key distribution. They are also commonly used in telecommunications for the connection between a source laser and a modulator, since the modulator requires polarized light as input. They are rarely used for long-distance transmission, because PM fiber is expensive and has higher attenuation than single-mode fiber.
PMMA [14-1]	Polymethyl Methacrylate – PMMA or poly(methyl 2-methylpropanoate) is the synthetic polymer of methyl methacrylate. This thermoplastic and transparent plastic is sold by the tradenames Plexiglas, Limacryl, R-Cast, Perspex, Plazcryl, Acrylex, Acrylite, Acrylplast, Altuglas, Polycast and Lucite and is commonly called acrylic glass or simply acrylic. The material was developed in 1928 in various laboratories and was brought to market in 1933 by Rohm and Haas Company.

POF [14-1]	Plastic Optical Fiber – POF is an optical fiber which is made out of plastic. traditionally PMMA (acrylic) is the core material, and fluorinated polymers are the cladding material. Since the late 1990s however, much higher-performance POF based on perfluorinated polymers (mainly polyperfluorobutenylvinylether) has begun to appear in the marketplace. In large-diameter fibers, 96% of the cross section is the core that allows the transmission of light. Similar to traditional glass fiber, POF transmits light (or data) through the core of the fiber. The core size of POF is in some cases 100 times larger than glass fiber. POF has been called the "consumer" optical fiber because the fiber and associated optical links, connectors, and installation are all inexpensive. The traditional PMMA fibers are commonly used for low-speed, short-distance applications in digital home appliances, home networks, industrial networks (PROFIBUS, PROFINET), and car networks (MOST). The perfluorinated polymer fibers are commonly used for much higher-speed applications such as data center wiring and building LAN wiring.
PSK [14-1]	Phase Shift Keying – A simple digital modulation technique that uses two different phase angles for 0 and 1.
Q	Quadrature – The quadrature component (90° out of phase) of a quadrature phase shift keyed signal and also the quadrature axis on a constellation diagram.
QPSK [14-1]	Quadrature Phase Shift Keying – A simple digital modulation technique that uses four phase shifts for each two bits of input.
QW [14-1]	Quantum Well – A QW is a potential well that confines particles, which were originally free to move in three dimensions, to two dimensions, forcing them to occupy a planar region. The effects of quantum confinement take place when the quantum well thickness becomes comparable at the de Broglie wavelength of the carriers (generally electrons and holes); leading to energy levels called "energy subbands", i.e., the carriers can only have discrete energy values.
REDFA	Rare Earth Doped Fiber Amplifier
RI [14-1]	Refractive Index – The RI (or index of refraction) of a medium is a measure for how much the speed of light (or other waves such as sound waves) is reduced inside the medium. For example, typical glass has a refractive index of 1.5, which means that light travels at 1/15 = 0.67 times the speed in air or vacuum. Two common properties of glass and other transparent materials are directly related to their refractive index. First, light rays change direction when they cross the interface from air to the material, an effect that is used in lenses and glasses. Second, light reflects partially from surfaces that have a refractive index different from that of their surroundings.

RIN [14-1]	Relative Intensity Noise – RIN describes the instability in the power level of a laser. The noise term is important to describe lasers used in fiber-optic communication and LIDAR remote sensing. Relative intensity noise can be generated from cavity vibration, fluctuations in the laser gain medium or simply from transferred intensity noise from a pump source. Since intensity noise typically is proportional to the intensity, the **relative** intensity noise is typically independent of laser power. RIN typically falls off with frequency and is a kind of pink noise. Relative intensity noise is measured by sampling the output current of a photodetector over time and transforming this data set into frequency with a fast Fourier transform. RIN is usually presented as relative noise power in decibels per hertz at one or several intensities.
ROADM [14-1]	Reconfigurable OADM – A ROADM is a form of optical add-drop multiplexer that adds the ability to remotely switch traffic from a WDM system at the wavelength layer. This allows individual wavelengths carrying data channels to be added and dropped from a transport fiber without the need to convert the signals on all of the WDM channels to electronic signals and back again to optical signals. The main advantages of the ROADM are: • The planning of entire bandwidth assignment need not be carried during initial deployment of a system. The configuration can be done as and when required. • ROADM allows for remote configuration and reconfiguration. • In ROADM, as it is not clear beforehand where a signal can be potentially routed, there is a necessity of power balancing of these signals. ROADMs allow for automatic power balancing. ROADM functionality originally appeared in long-haul DWDM equipment, but by 2005, it began to appear in metro optical systems because of the need to build out major metropolitan networks to deal with the traffic driven by the increasing demand for packet-based services.
RoF [14-4]	Radio-Over-Fiber – RoF refers to a technology whereby light is modulated by a radio signal and transmitted over an optical fiber link to facilitate wireless access. Although radio transmission over fiber is used for multiple purposes, such as in cable television (CATV) networks and in satellite base stations, the term RoF is usually applied when this is done for wireless access. In RoF systems, wireless signals are transported in optical form between a central station and a set of base stations before being radiated through the air. Each base station is adapted to communicate over a radio link with at least one user's mobile station located within the radio range of said base station.

RZ [14-1]	Return to Zero – RZ describes a line code used in telecommunications signals in which the signal drops (returns) to zero between each pulse. This takes place even if a number of consecutive zeros or ones occur in the signal. The signal is self-clocking. This means that a separate clock does not need to be sent alongside the signal, but suffers from using twice the bandwidth to achieve the same data-rate as compared to non-return-to-zero format. The "zero" between each bit is a neutral or rest condition, such as a zero amplitude in pulse amplitude modulation (PAM), zero phase shift in phase-shift keying (PSK), or mid-frequency in frequency-shift keying (FSK). That "zero" condition is typically halfway between the significant condition representing a 1 bit and the other significant condition representing a 0 bit. Although RZ contains a provision for synchronization, it still has a DC component resulting in "baseline wander" during long strings of 0 or 1 bits, just like the line code Non-return-to-zero.
SAN [14-1]	Storage-Area Network – A network of storage disks. In large enterprises, a SAN connects multiple servers to a centralized pool of disk storage. Compared to managing hundreds of servers, each with their own disks, SANs improve system administration. By treating all the company's storage as a single resource, disk maintenance and routine backups are easier to schedule and control. In some SANs, the disks themselves can copy data to other disks for backup without any processing overhead at the host computers.
SBS [14-1]	Stimulated Brillouin Scattering – Brillouin scattering occurs when light in a medium (such as water or a crystal) interacts with density variations and changes its path. The density variations may be due to acoustic modes, such as phonons, or temperature gradients. As described in classical physics, when the medium is compressed its index of refraction changes and the light's path necessarily bends.
SEED	Self-Electro-optic Effect Device
SFF	Small Form Factor
SG-DBR	Sampled Grating-Distributed Bragg Reflector
SiO_2 [14-1]	Silica or Silicon Dioxide – a hard glossy mineral which occurs naturally as quartz and is used in the manufacture of glass.
SM	Single-Mode
SMF [14-1]	Sing-Mode Fiber – In fiber-optic communication, an SMF (or SM fiber) is an optical fiber designed to carry only a single ray of light (mode). This ray of light often contains a variety of different wavelengths. Although the ray travels parallel to the length of the fiber, it is often called the transverse mode since its electromagnetic vibrations occur perpendicular (transverse) to the length of the fiber. SMFs are also called monomode optical fibers, single-mode optical waveguides, or unimode fibers.

SOA [14-1]	Semiconductor Optical Amplifier – SOAs are amplifiers which use a semiconductor to provide the gain medium. These amplifiers have a similar structure to Fabry-Perot laser diodes but with anti-reflection design elements at the endfaces. Recent designs include anti-reflective coatings and tilted waveguide and window regions which can reduce endface reflection to less than 0.001%. Since this creates a loss of power from the cavity which is greater than the gain it prevents the amplifier from acting as a laser. Semiconductor optical amplifiers are typically made from group III-V compound semiconductors such as GaAs/AlGaAs, InP/InGaAs, InP/InGaAsP and InP/InAlGaAs, though any direct band gap semiconductors such as II-VI could conceivably be used. Such amplifiers are often used in telecommunication systems in the form of fiber-pigtailed components, operating at signal wavelengths between 0.85 µm and 1.6 µm and generating gains of up to 30 dB. The semiconductor optical amplifier is of small size and electrically pumped. It can be potentially less expensive than the EDFA and can be integrated with semiconductor lasers, modulators, etc. However, the performance is still not comparable with the EDFA. The SOA has higher noise, lower gain, moderate polarization dependence and high nonlinearity with fast transient time. This originates from the short nanosecond or less upper state lifetime, so that the gain reacts rapidly to changes of pump or signal power and the changes of gain also cause phase changes which can distort the signals. This nonlinearity presents the most severe problem for optical communication applications. However it provides the possibility for gain in different wavelength regions from the EDFA. "Linear optical amplifiers" using gain-clamping techniques have been developed.
SOH	Silicon-Organic Hybrid
SP [14-1]	Service Provider – An organization that provides some kind of communications service, storage service or processing service or any combination of the three. Examples are a local or long distance telephone company, Internet service provider (ISP), application service provider (ASP) and storage service provider (SSP).
SPM [14-1]	Self-Phase Modulation – SPM is a nonlinear optical effect of light-matter interaction. An ultrashort pulse of light, when travelling in a medium, will induce a varying refractive index of the medium due to the optical Kerr effect. This variation in refractive index will produce a phase shift in the pulse, leading to a change of the pulse's frequency spectrum. Self-phase modulation is an important effect in optical systems that use short, intense pulses of light, such as lasers and optical fiber communications systems.

SRS [14-1]	Stimulated Raman Scattering – When a lower frequency 'signal' photon induces the inelastic scattering of a higher-frequency 'pump' photon in an optical medium in the nonlinear regime. As a result of this, another 'signal' photon is produced, with the surplus energy resonantly passed to the vibrational states of the medium. This process, as with other stimulated emission processes, allows all-optical amplification. Optical fiber is used as the nonlinear medium for SRS, for telecom purposes; in this case it is characterized by a resonance frequency downshift of ~11 THz (corresponding to a wavelength shift at ~1550 nm of ~90 nm). The SRS amplification process can be readily cascaded, thus accessing essentially any wavelength in the fiber low-loss guiding windows (both 1300 and 1550).
S-SEED	Symmetric SEED
TAM	Total Area Market
TDM [14-1]	Time Division Multiplexing – TDM is a type of digital or (rarely) analog multiplexing in which two or more signals or bit streams are transferred apparently simultaneously as sub-channels in one communication channel, but physically are taking turns on the channel. The time domain is divided into several recurrent timeslots of fixed length, one for each sub-channel. A sample, byte or data block of sub-channel 1 is transmitted during timeslot 1, sub-channel 2 during timeslot 2, etc. One TDM frame consists of one timeslot per sub-channel. After the last sub-channel the cycle starts all over again with a new frame, starting with the second sample, byte or data block from sub-channel 1, etc.
TEC [14-1]	Thermoelectric Cooler – A TEC is a solid-state active heat pump which transfers heat from one side of the device to the other side against the temperature gradient (from cold to hot), with consumption of electrical energy. Because heating can be achieved more easily and economically by many other methods, Peltier devices are mostly used for cooling. However, when a single device is to be used for both heating and cooling, a Peltier device may be desirable. Simply connecting it to a DC voltage will cause one side to cool, while the other side warms. The effectiveness of the pump at moving the heat away from the cold side is totally dependent upon the amount of current provided and how well the heat from the hot side can be removed.
TIA [14-1]	Telecommunications Industry Association – The TIA is a trade association in the US that represents about 600 telecommunications companies.

	Transimpedance Amplifier – The active current-to-voltage converter is an amplifier with current input and voltage output. The gain of this amplifier is represented by the resistance R ($K = V_{OUT}/I_{IN} = R$); it is expressed in units of ohms. That is why this circuit is named **transresistance amplifier** or more generally, **transimpedance amplifier**. Both terms are used to designate the circuit considered. Its input ideally has low impedance, and the input signal is a current. Its output may have low impedance, or in high-frequency applications, may be matched to a driven transmission line; the output signal is measured as a voltage.
TO [14-1]	Transister Outline metal-can package
TOSW	Thermo-Optic Switches
TPA [14-5]	Two-Photon Absorption – TPA is the simultaneous absorption of two photons of identical or different frequencies in order to excite a molecule from one state (usually the ground state) to a higher energy electronic state. The energy difference between the involved lower and upper states of the molecule is equal to the sum of the energies of the two photons. Two-photon absorption is many orders of magnitude weaker than linear absorption and is therefore not an everyday phenomenon. It differs from linear absorption in that the strength of absorption depends on the square of the light intensity, thus it is a nonlinear optical process.
TPPV	Two-Photon Photovoltaic
UNI	Ultrafast Nonlinear Interferometer
VCSEL [14-1]	Vertical-Cavity Surface-Emitting Laser – A type of laser diode that emits light from its surface rather than its edge. A VCSEL's circular beam is easy to couple with a fiber, and due to its surface-emission architecture, can be tested on the wafer. VCSELs are also noted for their excellent power efficiency and durability.
VCSOA	Vertical Cavity Semiconductor Optical Amplifier (SOA)

WDM [14-1]	Wavelength Division Multiplexing – In fiber-optic communications, wavelength-division multiplexing (WDM) is a technology which multiplexes multiple optical carrier signals on a single optical fiber by using different wavelengths (colors) of laser light to carry different signals. This allows for a multiplication in capacity, in addition to enabling bidirectional communications over one strand of fiber. "The true potential of optical fiber is fully exploited when multiple beams of light at different frequencies are transmitted on the same fiber. This is a form of frequency division multiplexing (FDM) but is commonly called wavelength division multiplexing." The term **wavelength-division multiplexing** is commonly applied to an optical carrier (which is typically described by its wavelength), whereas frequency-division multiplexing typically applies to a radio carrier (which is more often described by frequency). However, since wavelength and frequency are inversely proportional, and since radio and light are both forms of electromagnetic radiation, the two terms are equivalent.
XGM	Cross-Gain Modulation
XPM [14-1]	Cross Phase Modulation – XPM is a technique for adding information to a light stream by modifying the phase of a coherent optical beam with another beam through interactions in an appropriate non-linear medium. This technique is applied to fiber optic communications.

14.1 References

[14-1] http://encyclopedia.thefreedictionary.com

[14-2] http://en.wikipedia.org/wiki/Fiber-optic_communication

[14-3] Dutton, H. J. R.: Understanding Optical Communications, IBM Redbook 1998. http://www.redbooks.ibm.com/redbooks/pdfs/sg245230.pdf

[14-4] http://en.wikipedia.org/wiki/Radio_over_Fiber

[14-5] http://en.wikipedia.org/wiki/Two_photon_absorption