

Introduction to Optical Networks

P. Michael Henderson
michael.henderson@cox.net

September 17, 2001

*Waves are approaching.
Try to capture them. You can't.
Now they're particles.*

*Steve Wilson
Photonics Haiku¹*

Optical networks involve a number of technologies, from the physics of light through protocols. In fact, there's so much technology that most people involved with optical networks only have a full understanding of the narrow area they work in. This makes life difficult for the beginner who is approaching optical networks for the first time. This paper attempts to address this issue by providing a high level overview of many of the technologies involved in optical networks.

As a favor to me, I'd appreciate if you would send me an e-mail (at the address above) with nothing but a subject line of "Optical Networks" so that I can determine how many people are reading this paper. If you'd like to make comments, corrections, or offer suggestions about the paper, that would be appreciated, of course, but is not required.

The paper begins by discussing the physics of light, including lasers, optical modulation, photodetectors, and optical amplification principles. It then moves to the network level, describing the present state of optical networks, so-called First Generation networks, including SONET/SDH framing. Next it progresses to Second Generation "all optical" networks, and concludes by speculating on the emerging area of Metropolitan Optical Networks.

This is an extremely wide spectrum of information – consequently, you may find the information contained here to be somewhat brief and lacking in fine detail. Please let me know if there are areas you'd like to see discussed in more detail in future revisions of the document.

Physical Principles of Light

Electrons around the nucleus of an atom can exist only in discrete energy levels. When an electron transitions from a higher energy level to a lower energy level, energy is emitted from the atom². If the difference in energy between the two energy states falls within a certain range, the energy emitted is in the form of a quantum, or photon, of light.

Likewise, in order for the electron to move from a lower energy level to a higher energy level, the atom must absorb energy. Again, if the difference between the energy levels falls within a certain range, a

¹ Published in the "Peregrinations" column of the magazine *Photonics Spectra*, October 2000. In the column, the poetry was called "phoku" instead of photonics haiku (note: haiku consists of three lines, the first being five syllables, the second being seven syllables, and the third, five syllables).

² Photons are actually absorbed and emitted by electrons.

photon of light can provide the correct amount of energy to cause the electron to transition to the higher energy level.




Figure 1: The Bohr model of the atom. Electrons can only exist in discrete energy levels around the nucleus. To transition the electron to a higher level, the atom must absorb energy. When the electron transitions to a lower energy level, energy is emitted. Our interest is when the energy difference equals a photon of light.

Let's assume that the atom shown in Figure 1 has differences in its energy levels equal to a photon of light. If the electron is initially at level 2, it can transition to level 1 by emitting a photon. To get to level 3 from level 2, however, it must absorb a photon of the right frequency.

Building on this basic idea, there are three concepts that are important to optical theory: spontaneous emission, absorption of a photon (discussed above) and stimulated emission. Let me discuss each in turn. See Figure 2.

When atoms are in an elevated energy state, they sometimes enter a lower energy state when an electron spontaneously transitions to a lower energy level, emitting a photon in the process. This process is entirely random but there are so many atoms in any macroscopic amount of material that spontaneous emission occurs fairly often when the atoms are in an elevated energy state.

Absorption occurs when the electron absorbs a photon and moves to a higher energy level.

Stimulated emission is a bit more complex. Although Einstein first described it in 1917, it was not really exploited until the 1950's. In stimulated emission, a photon can interact with an atom and not be absorbed. If the atom is in an elevated energy state, the photon can cause the electron to transition to a lower energy state by emitting a photon which is exactly the same frequency as the first photon and in phase with it. This effect is exploited to build lasers.




Figure 2: Three important effects in optical physics.

Laser Operation

*Inside the laser
All the atoms march alone
But fall together*
Tony Harker
Photonics Haiku

A laser is built of a material which has electron energy band gaps equal to the energy of the photons desired (the frequency of the laser light times Planck's constant). For example, the first laser was built from synthetic ruby.

The laser has a mirror on each end of the laser material. One mirror is a full reflecting mirror while the other is a partial mirror. A partial mirror has the characteristic that photons have a probability of passing thru as well as being reflected. For example, a partial mirror might allow 70% of the photons to pass through while about 30% are reflected (some are absorbed and lost to the laser process).

Laser operation begins by injecting energy into the laser material in order to raise the atoms to a higher energy state. Atoms achieve a higher energy state when they absorb a photon and an electron transitions to a higher energy level around the nucleus. For continuous lasers, energy is added to the laser material continuously so that atoms which fall to a lower energy state are quickly pumped to the higher energy state. See Figure 3.

Some of the atoms in the elevated energy state transition to a lower energy state by emitting a photon as an electron spontaneously transitions to a lower energy level. If the photon is emitted in any direction except directly towards one of the mirrors, the photon is lost to the laser process. Some of the photons, however, strike a mirror at the proper angle and begin bouncing back and forth between the mirrors.

Invention of the Laser

Charles Townes and Arthur Schawlow filed for a patent on the laser in 1958, receiving it in 1960 (#2,929,922), but never built a working laser. Theodore Maiman of Hughes Research Labs built the first working laser in 1960, using synthetic ruby.

Gordon Gould, however, had been working on the laser in parallel with Townes and Schawlow and, after initial rejection by the patent office, was awarded a patent on the laser in 1977 (#4,053,845). Gould's patent was issued at a time when lasers were beginning to be used in great numbers, allowing Gould to collect significantly more in royalties than Townes and Schawlow ever received.

As they bounce back and forth, they interact with atoms which they pass and cause these atoms to emit photons at exactly the same frequency and phase. Based on the distance between the mirrors, one frequency of light becomes dominant and grows stronger. When the photons strike the partial mirror, some go through and this becomes the laser output.




Figure 3: Laser operation. Energy is put into the laser material to raise the atoms to a higher energy state. Some atoms spontaneously emit photons which bounce back and forth stimulating other atoms to emit photons. Some of the photons pass through the partial mirror, producing the laser output.

Laser Modulation

In order to communicate information using the laser, we need to modulate the light output. Although many kinds of modulation are possible, essentially the only technique used in the network today is called on-off keying (OOK), where the light is turned on or off to indicate a binary one or zero (see Appendix A for a more detailed discussion of OOK)³.

There are two common ways to modulate a laser: (1) direct modulation by controlling the drive signal to the laser, and (2) external modulation by placing a device which acts like a “shutter” between the laser and the fiber. Direct modulation is less expensive but generally causes a problem known as “chirp”. When the laser drive signal changes, the laser tends to change frequency slightly. This causes a problem because different frequencies of light propagate at different speeds in

Visible Light

The visible light spectrum extends from about 750 nm for deep red, to about 380 nm for deep violet. The 1310 and 1550 nm wavelengths used for communications, therefore, are in the infrared spectrum, invisible to the human eye.

³ Note that OOK of light is equivalent to amplitude modulation since light is an electromagnetic wave. There are two symbols per hertz of modulated signal but AM creates two side bands, each with a bandwidth of the modulated signal. Therefore, for OOK, the minimum bandwidth required (in hertz) is equal to the bit rate (10 Gbps requires a minimum of 10 GHz of bandwidth).

fiber (the lower frequencies travel slower than the higher frequencies). This is known as chromatic dispersion.

This variation in speed based on the frequency of the light causes the pulse to spread out as it propagates along the fiber. Eventually, the pulses merge together, causing a problem known as intersymbol interference (ISI), which makes it difficult to determine what the original pulse really was.

To mitigate this problem, external modulators, usually made of lithium niobate, are used. When an external modulator is used, the laser produces a constant output signal, which does not vary in frequency. The external modulator either blocks, or passes, the light in response to the electrical drive signal.

While the laser signal is constant in frequency, this does not eliminate chromatic dispersion. A pulse is actually made up of a number of frequencies which then travel at different speeds, again causing pulse spreading. It just happens less with a constant frequency laser source.

The problem of intersymbol interference is quite serious at higher signaling rates. As we operate at higher rates the symbol interval gets smaller and smaller. When we use OOK, each bit occupies about 100 ps when we operate at 10 Gbps and only about 25 ps when we operate at 40 Gbps. Almost all of the intersymbol interference is caused by dispersion. Even if dispersion was constant for each rate (which it is not) a fixed amount of dispersion will be much more serious as the symbol time gets smaller.

For example, suppose you had a fixed 10 ps of dispersion, independent of data rate. This would only be 10% of the symbol time at 10 Gbps but would be 40% of the symbol time at 40 Gbps.

One way to address this problem is to build better fiber, which has less dispersion. Another way is to put dispersion compensation elements in the network.

Another technique, which has been used in the electrical domain for a long time, is to use a modulation technique which carries multiple bits per symbol time. Since the dispersion problem is related to the symbol time, sending fewer symbols per second results in a longer symbol time. For example, if two bits could be sent per symbol, the symbol rate would be cut in half (each symbol would be twice as long) compared to OOK at the same bit rate.

A simple way to achieve this is to use multi-level optical signaling. For example, if four optical power levels can be used, two bits can be sent each symbol time. See Figure 4. This is easier said than done, however. First of all, lasers are not linear over the entire operational range used for OOK. When multi-level signaling is done, the laser must be limited to its linear range⁴. For a given laser, this reduces the amount of power launched into the fiber, but one can argue that a laser can be selected which has a linear operating power range which is equal to the non-linear operating power range of lasers usually used for OOK.

⁴ However, lasers are inherently non-linear and change their characteristics with temperature and with age.




Figure 4: Multi-level modulation of an optical signal. Instead of just turning the light off and on, the light is turned on at four different power levels, allowing two bits to be sent each symbol time. Note that the power is not zero for the lowest power symbol because we must stay in the linear portion of the laser characteristic curve.

At the receiver, the signal must be sampled and digitized to some number of bits. If everything were perfect, a four level signal could be detected with a device which sampled the signal and digitized it to two bits per sample. See Figure 5. But the world isn't perfect. For a two bit decoder to work, the signal must arrive at the detector such that the highest power transmit level is exactly equal to the full range of the detector.




Figure 5: Detection and digitization of the transmitted multi-level signal. This diagram is idealized in that the received signal is full scale of the detector.

But this is impossible in the real world. Different fiber links will have different amounts of attenuation and the power of the laser is likely to change as it ages. It's just not possible to guarantee that the maximum optical signal power received will be equal to the full range of the detector. So digitizing at two bits per symbol just won't work. Let's look at the case where the maximum power symbol received is equal to half the full range of the detector and the signal is digitized at three bits per symbol. See Figure 6.




Figure 6: Detection of a signal which is attenuated by 50% compared to the full range of the detector. Digitizing with three bits per symbol will work here.

The extra bit of accuracy in the digitizer provides dynamic range to the detector – the detector will work with signals which are “strong” as well as “somewhat weak”. But suppose the signal is attenuated by another 50%. Will three bits still work? See Figure 7.




Figure 7: Detection of a signal which is attenuated by 75% compared to the full range of the detector. Digitizing with three bits per symbol will not work here.

Here, we see that the signals fall between only two digitization levels, making it impossible to discriminate the four levels transmitted. So when the signal is only 25% of the full range of the detector, digitization at three bits will not work for a four level signal. Actually, for any signal less than about 37% of the full range of the detector, we no longer have four digitization levels and the detection fails.

Things are actually worse than shown here because the transmitted signal is not a nice square wave as shown in the diagrams and there is jitter in the signal, making it difficult to determine exactly when to sample the signal. All of these factors, plus others such as the linearity of the laser, go into the determination of the number of bits needed to digitize the signal.

Let's say that we analyzed the situation and determined that we can perform adequately with six bits per sample. The problem we then face is in building an analog to digital converter which can sample with six bits of accuracy 5 billion times per second (for 10 Gbps operation). This is not a trivial problem but one which will certainly be achieved in the future.

Photodetectors

Photodetectors are made from semiconductor material. It's important that the material be a conductor because our goal is to produce an electrical current from the optical signal. In this section we'll examine how this occurs.

The electrons in the semiconductor atoms can exist in a valance band or in a conduction band. Electrons in the valance band are tightly bound to the nucleus of the atom and cannot participate in current flow in the material. If the electron can transition to the conduction band, however, it becomes weakly bound to the nucleus and can participate in current flow in the material.

And, of course, the way that an electron transitions to a higher energy band is by absorbing a quantum of energy. If that quantum is equal to a photon of light at the frequency we're interested in, we can construct a photodetector. Figure 8 illustrates this. The valance and conduction bands are diffuse because the atom is in a crystal structure. The minimum energy gap between the valance and conduction bands is indicated by E_g . If the electron absorbs a photon with sufficient energy, it will transition to the conduction band.




Figure 8: The basic principle of photodetection in a semiconductor material. A photon is absorbed by the electron, causing the electron to transition to a higher energy band, the conduction band, where it becomes available for current flow. (source: adapted from [Ram98], p145).

A problem with constructing photodetectors for the 1310 and 1550 nm bands primarily used for long distance communications is that longer wavelength photons have less energy. Thus, photodetectors for use in these bands must be constructed of materials with relatively low energy band gaps. The most common materials are InGaAs and InGaAsP semiconductors.

To “collect” the electrons which move to the conduction band, we need to impose an electric field across the material so that the electrons in the conduction band move under the effect of the field. One way to achieve this is to form a junction using oppositely doped materials. Simply creating a junction will cause an electric field to form across the junction (called a depletion region)⁵ but a larger depletion region can be created by reverse biasing the diode. The larger depletion region improves the efficiency of the photodiode, as these types of photodetectors are called.

⁵ This effect was discovered early in the development of the semiconductor diode and transistor. Note that all diodes and transistors are enclosed in some type of opaque material (like the metal “top hat” cans of early transistors).




Figure 9: A semiconductor photodiode showing the enhanced depletion region due to the applied voltage. (source: adapted from a figure in [Ram98])

A problem with the photodiode in Figure 9 is that some of the semiconductor material is not part of the depletion region. Since this part of the material does not contain a strong electric field, electrons moved to the conduction band by absorption of a photon may not participate in the current flow. This causes a loss of efficiency in the photodiode.

To improve the efficiency of photodiodes, a lightly doped *intrinsic* semiconductor material is placed between the p-type and n-type materials. These photodiodes are called *pin* diodes because the material is made up of **p**-type, **i**ntrinsic, and **n**-type material. If the p-type and n-type material is chosen to be transparent to the frequency of light of interest, capture of photons occurs only in the intrinsic region, which experiences an electric field between the p-type and n-type materials. See Figure 10.

Pin diodes are often used for very high-speed signals, with an optical preamplifier.




Figure 10: A pin diode showing the p- and n-type materials and the intrinsic material between the doped materials.

The final type of photodiode we're going to discuss is the avalanche photodiode (APD). See Figure 11. The ADP, in a certain way, is a combination of the previous two photodetectors. It has an intrinsic region where the photons are captured, but it also has a junction (depletion layer). As with the pin diode, the materials are chosen so that all of the material except the intrinsic region is transparent to light at the wavelength we're interested in. Because of this, all photon capture occurs in the intrinsic region.




Figure 11: An avalanche photodiode (APD).

The major difference between the pin diode and the APD is the extra layer of p-type material between the n-type layer and the intrinsic layer. The junction of the p-type material and the n-type material form a junction which is then reverse biased by the voltage applied to the APD. This voltage is quite large, usually about 50 volts. Most of the voltage drop is across the junction between the n-type material and the p-type material, leading to a very strong electric field because the junction is not very wide.

When an electron is promoted to the conduction band in the intrinsic layer, it begins to migrate towards the junction due to the electric field across the whole photodiode. When it reaches the region of the junction it begins to experience a much stronger electric field and is accelerated. The electron gains significant kinetic energy through this increased velocity. If the electron collides with an atom, this kinetic energy can be transferred to the atom, allowing an electron in that atom to be promoted to the conduction band. That electron will see a strong electric field and will be accelerated, and perhaps to collide with another atom.

The original electron is not captured by the atom but bounces off after transferring its kinetic energy. It then accelerates again and can participate in another collision. In this way, the one electron promoted to the conduction band by the photon can be multiplied many times, leading to a much larger electric current for a given amount of light.

Unfortunately, APDs are not without their problems. First, the amount of amplification is limited. If too much amplification is provided, the device will “run away” and simply provide an essentially continuous (large) current. Second, APDs are noisy because thermal action can randomly promote an electron to the conduction band. The current resulting from this thermal action is noise. And finally, the avalanche action takes time to occur, meaning that the device has a certain response time. APDs have been used a great deal in the network but at the highest speeds, the equipment providers tend to use pin diodes with optical preamplifiers.

Optical Transmission and Amplification

Ordinary window glass has a high level of optical attenuation, due primarily to metallic ion impurities, especially iron⁶, but also copper, vanadium, and chromium. To use glass for communications, essentially all the impurities must be eliminated from the manufacturing process – a very difficult task.

The first practical glass fiber for telecommunications was produced by Robert Maurer, Donald Keck, and Peter Schultz of Corning in 1970. It was a single mode fiber with 17 dB/km of attenuation at the 633 nm window of the helium-neon laser. Later in 1970, Felix Kapron and his team at Corning developed a fiber optic cable with less than 20 dB/km loss in the 850 nm band of the GaAs/Al_xGa_{1-x}As semiconductor heterojunction laser⁷. By 1976, fiber loss had been reduced to less than 1 dB/km at 1310 nm, paving the way for the first commercial installation in Chicago in 1977⁸. See Figure 12 which contrasts the loss characteristics of early fiber with modern low loss fiber. The figure also shows the relationship of the windows used for communications over fiber with the visible light spectrum.

Some people question whether we should call the electromagnetic radiation used in fiber “light” since it is outside the visible band, but the term is convenient to use⁹. Be aware, however, that the frequencies involved force us to consider the quantum aspects of the radiation. We talk of the wave nature of the “light” when we consider the frequency spacing of multiple “beams of light” on a fiber, but we must consider light’s quantum (or particle) nature when we analyze the photodetection process, or how an atom emits a photon, as occurs in a laser.

⁶ The greenish cast to window glass is due primarily to iron impurities.

⁷ The first continuous wave, room temperature semiconductor laser was demonstrated in May 1970 by Zhores Alferov's group at the Ioffe Physical Institute in Leningrad (now St. Petersburg) and on June 1, 1970 by Mort Panish and Izuo Hayashi at Bell Labs. Zhores Alferov was awarded a Nobel Prize for his work on heterojunction lasers.

⁸ Operating at (only) 45 Mbps.

⁹ You’ll also hear people talking about “colors” of light, when the frequencies involved are not within the visible light spectrum.




Figure 12: The loss characteristics of early 1970's fiber optic cable and modern low loss fiber. Note, also, that the frequency bands used for communication are outside the visible spectrum, in the infrared. (source: adapted from a figure in [Bre99]).

Loss in glass optical fiber is due to either absorption or scattering and is dependent upon the wavelength of light. Pure silica exhibits significant Rayleigh scattering in the ultraviolet region above 400 nm¹⁰ and absorption in the infrared region below 2000 nm. Minimum loss (of about 0.2 dB/km) is in the 1550 nm region and is due primarily to Rayleigh scattering¹¹. See Figure 13 which details the loss profile of typical modern optical glass fiber. Note the significant peak of loss at about 1390 nm in modern fiber. This is due to absorption by the OH⁻ radical (from residual moisture in the fiber) which has a fundamental vibrational absorption peak at about 2730 nm. The overtones of this OH⁻ absorption peak are responsible for the dominant peak near 1390 nm and the smaller peak near 1230 nm. Another small peak occurs around 950 nm.

This absorption peak around 1390 nm does not preclude its use for communications – it simply adds attenuation in a non-linear manner across the band of about 1350 nm to about 1420 nm, which makes its use difficult. However, fiber optic manufacturing techniques have recently done much to reduce this OH⁻ attenuation peak¹² which will allow this band to be utilized for communications, further increasing the capacity of optical fiber. Soon, the entire band from about 1260 nm to 1625 nm will be able to be used

¹⁰ Rayleigh loss is about 30 dB/km at 400 nm. At 200 nm, the loss is in excess of 80 dB/km. The index of refraction peaks at the material's electron resonant frequency. Near this resonance, light is strongly absorbed. For glass, this occurs at wavelengths around 200 nm, resulting in the "UV protection" of glass.

¹¹ Rayleigh scattering is caused by random density fluctuations in the glass possibly created by thermal motion of the atoms as the glass cools. At some point in the cooling, the fluctuations are "locked in" to the glass.

¹² Lucent AllWave fiber claims a maximum of 0.31 dB/km across the 1400 nm band – less than the fiber's attenuation at 1310 nm, which can be as high as 0.39 dB/km. The zero dispersion point is about 1310 nm.

for communications, a bandwidth of over 50 terahertz. This could be especially useful in metropolitan areas where amplification is not required.




Figure 13: Measured loss profile (solid line) of a modern single mode fiber. The intrinsic loss (dashed curve) is due to Rayleigh scattering for wavelengths less than 1600 nm and absorption for wavelengths greater than 1600 nm (diagram taken from [Has95], p8).

Optical fiber consists of a central core, surrounded by a cladding of glass with a slightly lower index of refraction ($n_1 > n_2$)¹³. An opaque cover, designed to prevent leakage of light into the cladding and core, surrounds both of these. See Figure 14.

¹³ The index of refraction is varied by doping the silica with oxides of germanium, boron, phosphorus and various halides. The difference in the index of refraction between the core and the cladding is quite small, about 0.4%.




Figure 14: Cross-section of optical fiber showing the core, the cladding, and the opaque cover.

The two major types of fibers are multimode fibers and single mode fibers. Multimode fibers have a larger core, usually either $62.5 \mu\text{m}$ or $50 \mu\text{m}$ in diameter. Single mode fiber has a core which is about $8 \mu\text{m}$ in diameter (the most popular is $8.3 \mu\text{m}$ in diameter). The diameter with the cladding is $125 \mu\text{m}$ for both¹⁴. When the protective opaque covering is added, the diameter of a single fiber strand is about $250 \mu\text{m}$. The larger core in multimode fiber makes it is easier to couple to the light source, which may be a light emitting diode (LED). Multimode fiber, however, has significantly higher loss (due to modal dispersion) than single mode fiber and is therefore only used for short distance communications such as within a building or on a corporate campus. All long distance communications utilizes single mode fiber and laser light sources.

In multimode fiber, the light is guided by the almost perfect reflection at the interface between the core and cladding. However, in single mode fiber this effect is not what “guides” the light. Rather, the light can be thought of as a propagating electromagnetic wave which, in the absence of any light guide, would propagate along a spherical wave front. In a single mode fiber, the electromagnetic wave propagates in both the core and the cladding, with an exchange of energy between the two. The difference in the index of refraction between the core and the cladding causes the wave to propagate along a plane wave front, instead of a spherical wave front. A full understanding probably requires a review of the mathematics, available in [Ram98].

Speed of Light

All wavelengths specified are the wavelength of the light in a vacuum ($c = 2.99792458 \times 10^8 \text{ m/s}$). The speed of light in fiber is about $0.66c$, which would make the wavelengths shorter if measured in fiber.

By happenstance, the speed of light in fiber is very close to the speed of electricity in copper ($\sim 0.68c$) so the propagation delay in an optical network is about the same as it was in the pre-optical network.

¹⁴ When Corning was developing their early optical fiber, they found that thicker fiber (they tried $250 \mu\text{m}$) was too brittle and tended to break, while smaller fiber ($50 \mu\text{m}$) tended to stick to the pulleys used to spool the fiber. They settled on $125 \mu\text{m}$ which has remained the standard today.

The bandwidth of fiber in the 1550 nm range is about 12 terahertz and in the 1310 nm range about 5 terahertz. With the elimination of the OH⁻ absorption peaks, the usable bandwidth will be over 50 terahertz. Assuming a signal to noise ratio of about 100, Shannon's theorem tells us that the ultimate capacity of the fiber in this 50 terahertz band is about 333 terabits per second¹⁵.

The characteristics of single mode fiber have developed and changed over time. Early fiber had zero dispersion¹⁶ in the 1310 nm region. Since dispersion is a problem in communications, later fiber was modified so that the zero dispersion point was shifted to the 1550 nm region. Japan has made a significant effort to install fiber and installed much of their fiber during this period (so much of the fiber in Japan is this dispersion shifted fiber).

When wavelength division multiplexing (WDM) came along, however, it was discovered that some (local) dispersion was necessary for proper operation of WDM¹⁷. Most modern fiber, therefore, is not dispersion shifted to zero dispersion in the 1520 to 1580 nm range.

For standard glass fiber, there's a lower limit of about 0.15 dB on the attenuation per km due to Rayleigh scattering. Thus, the signal must be amplified for long distance transmission (see Figure 15). In the early days of optical communications, this was done electronically by detecting the signal and retiming, reshaping, and retransmitting it (known as 3R regeneration). Electronic regeneration has a number of problems, however, including equipment cost and the problem of locating and powering the regeneration equipment. For example, if a fiber were carrying even 32 wavelengths, every regeneration point would have to consist of 32 detectors, 32 lasers, plus electronics and the equipment to separate and join the different wavelengths. This would be prohibitively expensive since the regeneration might have to be done every 40 km.

Dispersion Compensation

Some fiber is dispersion shifted to zero below 1520 nm, while other fiber is shifted to zero dispersion above 1580 nm (note that this second fiber has negative dispersion across the C-band). The two types of fiber are combined on a route to provide as close to zero dispersion across the entire route as possible. Thus, while zero dispersion is undesirable on a single link, it is desired across the entire route.

Alternately, dispersion compensation elements are used at EDFA points instead of using different fibers over a route.

¹⁵ Shannon's theorem is $C = B \log_2 (1 + S/N)$, where C is the channel capacity in bps, B is the bandwidth in hertz, S is the signal power, and N is the noise power. Coming close to this capacity will require line coding which produces a spectral efficiency of 6 bits/symbol/Hz (the Hz refers to bandwidth). Present commercial technology (on-off keying) can theoretically achieve 1 bit/symbol/Hz but the limits of real-world filters mean we achieve less.

¹⁶ Dispersion is caused by different components of the light traveling at different rates in the fiber. Ordinary fiber has zero dispersion around 1310 nm.

¹⁷ With no dispersion, the signals in the different λ s (lambdas or wavelengths) are phase matched and generate four-wave mixing, which degrades performance. With some dispersion, the signals in the different λ s do not remain in phase.




Figure 15: Even though fiber has low loss, amplifiers are required to regenerate the optical signal when transmitted over long distances.

The problem of amplifying optical signals for long distance transmission was successfully addressed by the development of Erbium doped fiber amplifiers (EDFAs). Erbium doped fiber amplifiers had their roots in the early (1964) amplification experiments in rare earth doped fiber lasers. Practical EDFAs, however, did not appear until about 1988¹⁸.

An EDFA consists of a length of silica fiber with the core doped with ionized atoms (Er^{3+}) of the rare earth element Erbium. This fiber is “pumped” with a laser at a wavelength of 980 nm or 1480 nm¹⁹. This doped, pumped fiber is optically coupled with the transmission fiber so that the input signal is combined with the pump signal in the doped fiber. An isolator is used at the input and/or output to prevent reflections which would convert the amplifier into a laser.

Erbium has multiple energy levels, three of which are of interest to this discussion, and are labeled E_1 , E_2 , and E_3 in Figure 16. Since the Erbium ions are part of the glass fiber, the energy bands are split into multiple energy levels, a process known as *Stark splitting*. The net effect is that at the macroscopic level, the Erbium ions appear to have a continuous energy band in the region of each discrete energy band shown in Figure 16.

When pumped by a 980 nm laser, ions that are normally at the rest state E_1 are pumped to state E_3 . This state has a short lifespan of about 1 μs , decaying to state E_2 which has a relatively long lifespan of about 10 ms. If the pump laser has sufficient power, the population of Erbium atoms will eventually stabilize around energy level E_2 .

A photon of the input signal will cause a pumped ion to transition from the E_2 state to the E_1 state, emitting a photon with the same energy (frequency or wavelength) as the incident photon. In the process, the emitted photon may cause other Erbium ions to emit photons with the same energy, providing amplification. Since amplification is by stimulated emission, i.e., the output wavelength is equal to the input wavelength, EDFAs work well for wavelength division multiplexer (WDM) systems which utilize a range of wavelengths.

¹⁸ S. B. Poole at the University of Southampton did some of the pioneering work on Erbium doped fiber amplifiers. David Payne and P. J. Mears, of the University of Southampton and Emmanuel Desurvire of Bell Labs were also instrumental in developing practical EDFAs.

¹⁹ Note that photon energy is inversely related to wavelength. The energy of a photon is equal to Planck’s constant (6.626×10^{-34} Joule-sec) times the frequency ($E = hf$ or $E = hc/\lambda$). A photon at 980 nm has more energy than one at 1480 nm, which has more energy than one at 1530 nm. Like most things in life, you put in more energy than you get out.




Figure 16: Three energy levels of Er³⁺ Erbium doped into silica fiber. Note that the energy levels are spread into bands (gray in the figure) due to the Stark splitting process. (source: adapted from [Ram98], p123)

Early EDFAs could provide 30 to 40 dB of gain in the C-band²⁰ of 1530 to 1565 nm with noise figures of less than 5 dB²¹. Recently, new EDFAs have been developed which can provide 25 dB of gain in the L-band (1565 to 1625 nm) as well as the C-band. Researchers are now attempting to develop amplifiers for the S-band (1460 nm to 1530 nm). Our ability to utilize even shorter wavelength bands for long distance transmission, all the way to perhaps 1260 nm, will depend on researchers' ability to develop optical amplifiers for these bands.

Lasers at 1480 nm can also be used for pumping but generally cannot achieve full population inversion, as the 980 nm lasers can. However, a 1480 nm signal can be transported through the fiber, allowing operation of EDFAs without electrical power.

EDFAs for the 1310 nm band are not available at this time although work is being done on praseodymium doped fiber amplifiers.

Recently, a new type of amplifier, known as a Raman amplifier, has been announced. Raman amplification exploits an effect known as stimulated Raman scattering (SRS). When two or more signals at different wavelengths are injected into a fiber, power is transferred from the shorter wavelength signal to the longer wavelength signal, known as the *Stokes wave*. The Stokes wave must be within 60 to 100 nm of the pump wave²². The effect operates in both directions. Thus, the pump signal could be propagating in one direction, while the information signal is propagating in the other and the energy exchange will take place. SRS also occurs when the signals are propagating in the same direction.

²⁰ The "C-band" is so named because it is the "conventional" band. The L-band is named for "long-wavelength". The S-band is named for "short-wavelength". Other bands are: the O-band for "original" from 1260 to 1360 nm, the E-band for "extended" from 1360 to 1460 nm, and the U-band for "ultra long wavelength" from 1625 to 1675 nm.

²¹ The noise figure, in dB, is equal to the input SNR in dB, minus the output SNR in dB. Even though you get amplification, the SNR degrades. One source of noise is spontaneous emission by the atoms in the EDFA.

²² According to [Ram98], peak coupling occurs between the pump signal and the Stokes wave when the Stokes wave is about 11 to 13 terahertz below the pump frequency. If the Stokes wave is to be at 1550 nm, the pump wave should be 83 to 98 nm above (shorter) for optimum coupling.

SRS occurs in standard fiber, not requiring any doping or special fiber characteristics. As such, Raman amplification can be added to existing installed fiber.

SRS is dependent upon the energy density in the fiber, being more pronounced as the energy density increases. With the small core radius of single mode fiber, it is fairly easy to reach energy densities where SRS is pronounced.

Rather than attempt to eliminate SRS, Raman Amplifiers exploit this effect. Network equipment companies have announced Raman amplifiers which launch the pump signal counter to the information signals to be amplified. This could be done from the location of an EDFA, for example. With the right choice of wavelength(s) for the pump signal and the information signals, power is transferred from the pump signal(s) to the information signals.

The gain provided by today's Raman amplifiers is modest and does not eliminate the need for EDFAs²³. Its greatest value may come as the network converts from OC-192 to OC-768 (10 Gbps to 40 Gbps). OC-768 would normally require closer spacing of the EDFAs. Since Raman amplification can be done from existing EDFA sites by inserting a pump λ (lambda or wavelength) at the proper place in the spectrum, it may allow existing fibers to be used for OC-768, without having to add additional (more closely spaced) EDFAs.

Additionally, Raman amplification can be done across the entire spectrum of interest, from 1260 nm to 1625 nm, a bandwidth of over 50 terahertz. When Raman amplifiers which can provide gain equivalent to EDFAs are paired with low OH⁻ fiber, they will open up the possibility of using the entire spectrum from 1260 nm to 1625 nm for long distance dense wavelength division multiplexing (DWDM) communications, significantly increasing the backbone capacity of optical networks.

Raman amplification is not free, however. At a minimum, it uses up λ s which could be used for information transfer. It only pays if the speed of the information carrying λ s increases more than the loss of the λ s used for Raman amplification, or the available bandwidth increases significantly. Thus, if Raman amplification allows a fiber optic span to be converted from OC-192 to OC-768, the capacity of the information carrying λ s increases by four. Even if one pump λ has to be used for each information carrying λ , there will be a net gain in the information carrying capacity of the fiber span.

First Generation Optical Networks

The basis of first generation networks is Synchronous Optical Network (SONET) and Synchronous Digital Hierarchy (SDH) techniques. Before describing the topology of first generation networks, I will first address SONET/SDH.

SONET defines the low level framing protocol used primarily on optical links. It was developed in the United States through the ANSI T1X1 committee. ANSI work commenced in 1985 with the CCITT (now ITU) initiating an effort in 1986. From the very beginning there was conflict between the US proposals

²³ The Raman amplifier can be viewed as the preamp in a cascaded amplifier system consisting of the Raman amplifier and the EDFA. Raman amplifiers can have better noise figures than EDFAs. Since the effective noise figure of a cascaded amplifier is highly dependent on the noise figure of the preamp, these systems can provide amplification with less SNR degradation than a system with EDFAs only. The equation for the noise figure of a cascaded amplifier is $F = F_1 + (F_2 - 1)/G_1$ where the subscripts refer to the gains stages. Note that this equation is not in dB.

and the ITU. The US chose a data rate close to 50 Mbps in order to carry their T1 (1.544 Mbps) and T3 (44.736 Mbps) signals. The European delegates needed a specification which would carry their E1 (2.048 Mbps) as well as E3 (34.368 Mbps) signals efficiently. The Europeans rejected the 50 Mbps proposal as bandwidth wasteful and demanded a base signal rate close to 150 Mbps²⁴.

Eventually a compromise was reached which allowed the US data rates to be a subset of the ITU specification, known formally as Synchronous Digital Hierarchy (SDH).

Since SONET is a subset of SDH, I'm going to begin my discussion with an explanation of the SONET frame, and then generalize to SDH.

The basic SONET frame is set up as shown below, in Figure 17, as 9 rows of 90 octets. It is transmitted from left to right and top to bottom. That is, the octet in the upper left corner is transmitted first followed by the second octet, first row, etc.

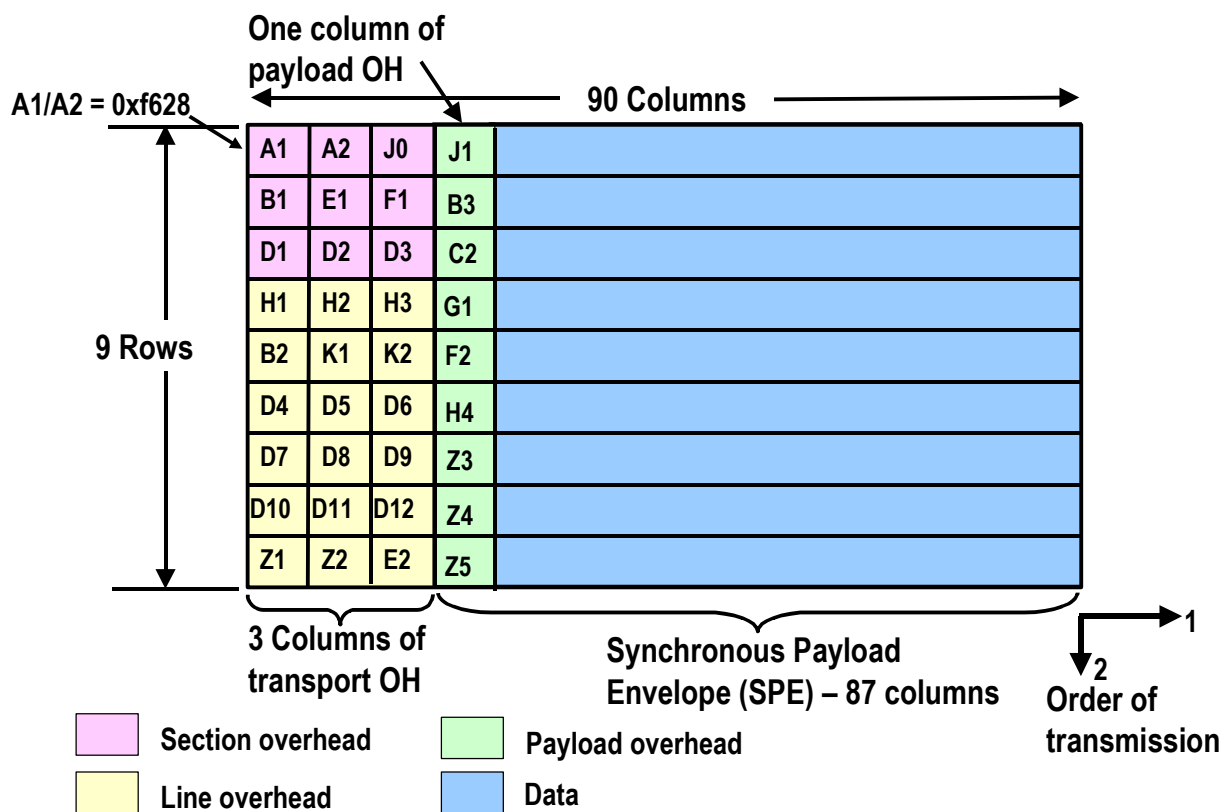


Figure 17: The SONET frame.

Framing is accomplished by the first two octets, called the A1 and A2 octets. When the frame is transmitted, all octets except A1, A2, and J0 are scrambled to avoid the possibility that octets in the frame might duplicate the A1/A2 octets and cause an error in framing. The bit pattern in the A1/A2 octets is

²⁴ There's good and bad in both proposals. SDH's granularity of 155 Mbps is too large for commercial services today, but adoption of SONET's 52 Mbps granularity causes excess overhead when higher data rates are utilized. SDH carries the entire SONET overhead in order to be compatible with it – an STM-1 carries 9 columns of overhead although only three columns are needed.

1111 0110 0010 1000 (0xf628). The receiver searches for this pattern in multiple consecutive frames²⁵, allowing the receiver to gain bit and octet synchronization (once bit synchronization is gained, everything is done, from there on, on octet boundaries – SONET/SDH is octet synchronous, not bit synchronous).

The first three columns of a SONET frame are called the Transport Overhead (TOH). The 87 columns following the TOH are called the Synchronous Payload Envelope (SPE). Within the SPE there is another column of overhead, called the Payload Overhead (POH), whose location varies because of timing differences between networks. I won't go into a great deal of detail about why the location of the POH varies – if you're interested, you should pick up one of the texts listed in the references. This leaves 86 columns by 9 rows for usable payload in an OC-1.

Remember that everything about SONET/SDH is tied to the transmission of G.711 PCM encoded voice. G.711 samples voice 8,000 times a second, or every 125 μ seconds, and, as might be expected, the SONET frame repeats 8,000 times per second.

This gives a data rate of 51.84 Mbps (90 columns times 9 rows, times 8,000 times per second, times 8 bits per octet). This signal is known as a Synchronous Transport Signal – Level 1 or STS-1. Once the scrambler is applied to the signal, it is known as an Optical Carrier – Level 1 signal or OC-1.

Since there are 86 non-overhead columns of 9 rows, a SONET OC-1 frame has a usable payload rate of 49.536 Mbps, sufficient bandwidth to carry 774 simultaneous voice conversations. This is in excess of the 672 simultaneous voice conversations carried in a T3, allowing T3s to be easily mapped into a SONET channel.

If every T1 and T3 carried nothing but voice signals, it would be possible to de-multiplex the T1/T3 signal and carry the voice channels in SONET native mode, allowing greater bandwidth efficiency. However, T1s and T3s are often used to carry data instead of multiple voice channels. The T1 or T3 must, therefore, be carried as a single data stream instead of multiple voice channels. This requires additional overhead in the SONET frame.

This plesiochronous digital hierarchy (PDH) traffic (T1, T3, E1, or E3) is encapsulated with additional framing octets, designed to allow the PDH traffic to be carried within a SONET/SDH channel. This is known as a virtual tributary (VT) in SONET and a virtual container (VC) in SDH. Multiple T1 circuits, for example, may be combined into a single SONET channel, up to 28 T1s in an OC-1.

I mentioned earlier that the ITU established a base rate close to 150 Mbps for SDH. Specifically, the rate the ITU established is three times the US OC-1 rate (155.52 Mbps) and is called the Synchronous Transport Module – Level 1 or STM-1. SDH also uses a 9 row frame but has three times as many columns as the OC-1 signal (270 octets instead of 90 octets). In order to make it easy to mux and demux the signals, everything is scaled by three from the OC-1 frame. Thus, in an SDH STM-1 there are 9 octets of Transport Overhead per row and three octets of Path Overhead. See Figure 18.

²⁵ The framing hardware searches for the boundary between A1 and A2 for octet synchronization. It must find the A1/A2 pattern for a certain number of consecutive frames before it leaves the seek state and enters the synchronized state.

A1	A1	A1	A2	A2	A2	J0	Z0	Z0	
B1	X	X	E1	X	X	F1	X	X	
D1	X	X	D2	X	X	D3	X	X	
H1	H1	H1	H2	H2	H2	H3	H3	H3	
B2	B2	B2	K1	X	X	K2	X	X	
D4	X	X	D5	X	X	D6	X	X	
D7	X	X	D8	X	X	D9	X	X	
D10	X	X	D11	X	X	D12	X	X	
S1	Z1	Z1	M0/1	Z2	M2	E2	X	X	

First STS-1
 Second STS-1
 Third STS-1

Figure 18: An STS-3 SONET frame or an STM-1 SDH frame. For SONET, the number of transport overhead columns will be equal to three times N where N is the value in STS-N. For SONET, N is defined as 1, 3, 12, 48, 192 and 768. For SDH, the number of overhead columns is nine times N where N is the value of the STM-N. Common STM-Ns are 1, 4, 16, 64 and 256. Other Ns are possible for both SONET and SDH but are not normally used.

When multiple channels of OC-1 are transmitted, the data is octet multiplexed. For example, an OC-3 signal will transmit octet A1 of stream 1, then octet A1 of stream 2, then octet A1 of stream 3, then octet A2 of stream 1, octet A2 of stream 2, etc. This multiplexing is carried out for all levels of SONET and SDH, including OC-192 and OC-768. Because of this, SONET/SDH maintains a frame time of 125 μ s.

Just as the OC-1 rate is the base of the US SONET, the STM-1 rate is the base of the SDH. Multiples of STM-1 are defined, all as multiples of 155.52 Mbps. See Table 1 below which summarizes the different rates for both SONET and SDH.

SONET name	SDH name	Line rate (Mbps)	Synchronous Payload Envelope rate (Mbps)	Transport Overhead rate ²⁶ (Mbps)
STS-1	None	51.84	50.112	1.728
STS-3	STM-1	155.52	150.336	5.184
STS-12	STM-4	622.08	601.344	20.736
STS-48	STM-16	2,488.32	2,405.376	84.672
STS-192	STM-64	9,953.28	9,621.504	331.776
STS-768	STM-256	39,813.12	38,486.016	1,327.104

Table 1: SONET/SDH digital hierarchy.

Sometimes a data stream is transmitted not as a sequence of voice channels, but as a single data stream. When this data rate exceeds an OC-1 or an STM-1 it is considered a concatenated, or unchannelized, data stream and is indicated with a lower case “c” following the name. Concatenated channels have their payloads locked, i.e., when slippage occurs it is done simultaneously for the entire frame and is for N octets. Thus, an OC-3c is a single payload (SPE) data stream of 150.336 Mbps. An OC-48c is a single payload (SPE) data stream of 2.405 Gbps.

There are many other aspects of SONET/SDH which are not covered here. Probably the most important subject not covered is how SONET/SDH handles clock differences between networks. While this is an interesting topic, it is not directly applicable to our discussion. Those wishing further information should consult the references, especially [Hen01].

SONET Networks

Optical networks can be configured point-to-point, linear, ring, or mesh. First generation networks, however, were deployed before techniques to manage complex mesh networks were developed. Consequently, first generation optical networks are primarily rings, although point-to-point and linear networks are utilized for certain applications.

A point-to-point network is simply that – a network of one optical link, with terminating multiplexers (TMs) at each end. A linear network is similar to a point-to-point network, but contains intermediate nodes, called add/drop multiplexers (ADMs). In practice, the TMs are the same as the ADMs but only have one optical connection. A ring is a linear network which folds back on itself. See Figure 19 which shows the layout of point-to-point, linear, and ring optical networks.

The defining characteristic of “first generation” optical networks is that the optical signal is converted to electronic at each ADM. Second generation optical networks do not convert the optical signal to electronic at the optical add/drop multiplexer. There is a bit of crossover, however. When wavelength division multiplexing came along, some first generation ADMs were modified to allow some wavelengths to pass through the ADM without being converted to electronic. These wavelengths were called “express lambdas,” named after express trains which do not stop at each station.

²⁶ Overhead associated with the transport overhead columns only. Excludes overhead contributed by the POH.

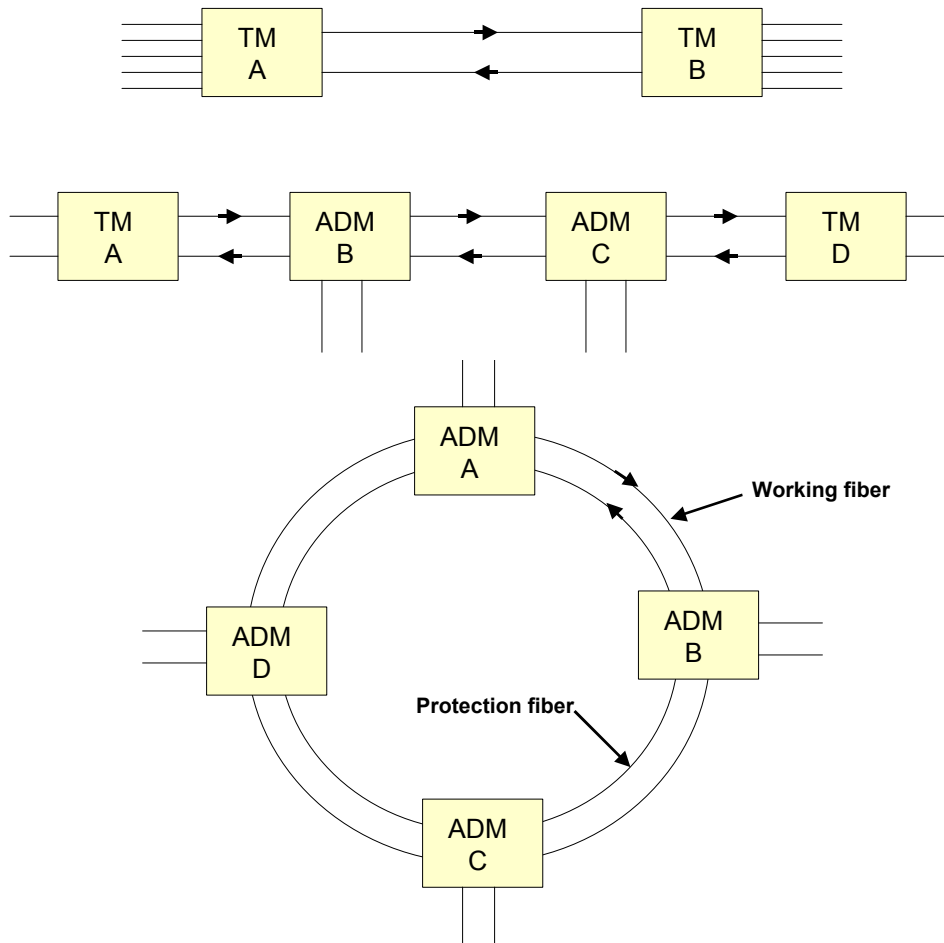


Figure 19: Point-to-point, linear and ring optical networks. A terminating multiplexer (TMs) is essentially the same as an add/drop multiplexer (ADM) but without the second optical interface.

ADMs are used to insert and remove SONET channels (like an OC-1 channel from an OC-48 fiber). A SONET channel may be carrying multiple T1 or T3 circuits in a virtual tributary (VT). Modern ADMs can extract or insert SONET channels or PDH (T1 or T3) channels.

Traditionally, PDH channels were dealt with by digital cross connect systems (DCSs). A DCS is essentially an electronic patch panel, where channels are interconnected. A DCS may interconnect DS0s of T1s, or T1s of T3s, or combinations. With the introduction of SONET, the DCSs evolved into combined DCS/ADM equipment.

By far, the most common topology for first generation networks is the ring because of the need to maintain service in the face of equipment failures and fiber cuts. There are two types of ring architectures used in first generation networks: *unidirectional path switched rings (UPSR)* and *bi-directional line switched rings (BLSR)*. Let me discuss the unidirectional/bi-directional part first.

Suppose we have a ring with four ADMs, labeled A, B, C, and D as shown in the ring in Figure 19. If data is to be sent from A to B, it is sent in the most direct way, across link A-B. But what is the return path, from B to A? With a unidirectional ring, all data is sent in the same direction, so the data going

from B to A traverses ADMs C and D before arriving at A. In a bi-directional ring, the data is sent in the opposite direction, directly from B to A.

Now, why do we have these two architectures and what are the advantages and disadvantages of each? Unidirectional rings are simpler to administer but the link delay is asymmetric. That is, the propagation delay from B to A is longer than the propagation delay from A to B. This effect is exacerbated when the ring is large, so we find unidirectional rings used mostly in metropolitan areas, while bi-directional rings are used for longer spans.

Protection, the term used for recovering from a fiber or electronics failure, can be performed in two ways in a ring. In path switching, the data sent between two ADMs is duplicated on two fibers, one transmitting clockwise and one counterclockwise.

The receiving ADM has duplicate electronics and receives and decodes the data from both fibers, selecting the data from the "working fiber" until a problem is detected. It then switches to the data in the "protection fiber," normally without causing any interruption to the traffic stream. Path switching is fairly simple since the receiver can both detect the problem and switch to the protection fiber (no coordination is necessary with the transmitter). No communication is required between ADMs to effect the switch. This type of protection is known as "1+1" protection and requires two times the fiber bandwidth. In general, any type of backup scheme which provides the ability to completely restore all services requires two times the fiber bandwidth.

Line switching on the other hand involves maintaining a backup fiber and switching to it when a problem is detected in the working fiber. Since the transmitter must do the switching but the receiver detects a cut fiber, this type of protection can be more complex. If the transmit and receive fibers are run in the same bundle, a backhoe is likely to cut both at the same time, allowing both ends to recognize the failure and do the switch. When the failure is due to the more common problem of electronics failure, the receiver must detect the failure and send a message to the transmitter to effect the switch. This takes more time but the restoration can be completed within the 50 ms requirement.

Line protection is normally done on a four-fiber ring, with two working fibers and two protection fibers. If only one fiber is cut, the circuit can be restored by switching to the protection fiber. This is known as "span protection." If all the fibers are cut, or if an ADM fails completely, the link is restored by routing the traffic in the other direction around the ring, and is known as "line protection." Line switching can be accomplished on a two-fiber ring by using only half the bandwidth of the fiber for primary traffic, reserving the other half for backup. In this case, span protection is not possible, but line protection works much the same as for the four-fiber system.

When a backup fiber is dedicated to backing up each working fiber, the protection is known as 1:1. It is possible, however, to have one fiber backup a number of fibers, known as 1:n protection. When 1:n protection is used, the traffic must be prioritized because some will be dropped if the number of failures exceeds the number of backup fibers.

When wavelength division multiplexing (WDM) is used, the same type of protection can be provided, although the physical topology must be recognized. For example, in a two fiber system, it is possible to

Spatial Reuse

When data is sent around a unidirectional SONET ring, the time slots are not reused by any ADMs beyond the destination ADM. This is known as "source stripping."

With bi-directional rings, however, it is possible to reuse the time slots beyond the destination ADM, a technique known as "destination stripping" or "spatial reuse." Managing a network with spatial reuse is more complex than one without it, however, and in certain topologies does not provide any value. Spatial reuse is primarily used in long distance networks.

reserve one or more wavelengths, or λ s, instead of reserving half of the bandwidth in a λ . Thus, path switching can be implemented with two λ s on two fibers, instead of using half the bandwidth on a single λ . When line switching is utilized, backup can be provided by utilizing available λ s on the protection fiber, instead of reserving half the bandwidth in a single λ . Since physical fibers get cut, and not λ s, the physical topology of the network must be understood when planning the protection.

Note that when WDM is used, all λ s are not terminated at each ADM. Depending upon the traffic, some λ s (sometimes called “express λ s”) can be routed through the ADM, remaining in optical form and reducing the cost of electronics in the ADM. In a way, this routing of λ s was the genesis of second-generation optical networks, which we discuss in a later section.

The above discussion focused mainly on the fiber failures. Although ADMs have significant redundant equipment (dual power supplies, redundant switch fabrics) catastrophic failures occur and must be planned for. One problem area is where networks interconnect. For example, a metropolitan network will connect to a long haul network at an ADM. If a single ADM is used, a failure can completely disrupt communications between the networks. To address this issue, the concept of “dual homing” is used, where two ADMs in each network, physical separated, are used to connect networks together (see Figure 20). These ADMs operate by dropping the traffic for the next network, but simultaneously passing it through to the other ADM. Thus, both ADMs are passing the same traffic to the other network, and the other network uses one link as a primary link and the other as a protection link. If one of the ADMs fails, the network switches to the protection link.




Figure 20: Interconnection of two SONET/SDH rings showing dual homing between the networks. Often, the interconnect is via digital cross connect systems (DCS), rather than ADMs.

ATM and IP transport in First Generation Optical Networks

First generation optical networks are essentially fixed bit pipes – provisioning a circuit takes a relatively long time, certainly when compared to the circuit set up time for traditional telephony. Additionally, the level of granularity in the bit rates is fairly large – 52 Mbps in SONET and 155 Mbps in SDH. Clearly, some higher level transport technology must be employed which can aggregate and rapidly route traffic. The two major contenders which have emerged are asynchronous transfer mode (ATM) and Internet protocol (IP).

ATM is a fixed size cell technique. It is a connection based system, where a route is established and all cells for that connection follow the same route. ATM was designed with a quality of service (QoS) feature primarily to allow it to handle voice traffic. The format of the ATM cell allows the receiver to find the beginning of the frame and to detect loss of frame sync.

IP on the other hand is a connectionless system, without QoS, and requires a lower level protocol for framing. The majority of traffic on the network today is IP.

ATM can transport IP traffic through use of the ATM adaptation layer 5 (AAL-5), which is essentially a protocol conversion technique. IP network designers, however, do not see the value add of using ATM and have designed techniques for transporting IP directly over SONET. The primary technique is packet over SONET (POS) which utilizes an octet oriented HDLC frame to transport the IP packets [Sim62].

POS utilizes a framing octet (0x7E) to indicate the start and end of the HDLC frame. A shielding octet (0x7D) is used to shield any octets in the IP frame that duplicate the framing octet or the shielding octet. POS is defined up to OC-48, and a variation is being developed for OC-192 and above (in one proposal, the framing and shielding become two character sequences) [Mca99].

The need to insert shielding characters causes the frame to be lengthened in a probabilistic manner (depending upon the occurrence of flag and shielding characters in the IP frame), making transmission of IP frames non-deterministic. To address this issue, Lucent has proposed an alternative technique for framing IP (actually, any data) over SONET called "simplified data link" (SDL) [Her99][Dos99]. SDL uses a framing technique similar to ATM, and does not require a framing symbol (bit, octet, or multiple octet). Because of this, frames are not expanded, as occurs with POS. SDL has been implemented by Lucent and PMC-Sierra in their POS chip sets, allowing their customers to choose between POS and SDL. POS is extremely well established, however, so we'll have to wait and see if SDL can gain a foothold in the market. The ANSI T1X1.5 committee has proposed a technique called Generic Framing Procedure (GFP) for framing variable length data. While not identical, GFP looks a lot like SDL.

The trend away from ATM for IP transport will likely continue. IP routers are adding QoS through multiprotocol label switching (MPLS) and the protocol conversion to and from AAL5 is slower and more difficult than POS. But perhaps the main reason for the migration away from ATM is that the ATM network is a separate layer that must be managed, separate and apart from the optical network and the IP network. Unless the ATM layer provides clear advantages, it will be bypassed (which is what seems to be happening in the backbone).

ATM certainly has a place in the access network, where it provides fine levels of bit rate granularity, QoS, and connections on demand. As corporations need higher access speeds, however, they may accept SONET's 52 Mbps granularity and IP's MPLS QoS.

Second Generation Optical Networks

The essence of second-generation optical networks is that traffic stays in optical form throughout the network, being converted to electronic only at the edges of the network. Second generation optical networks assume WDM and the λ becomes the lowest level of transmission granularity.

The transition to an all-optical network must be evolutionary, meaning that second generation networks must carry first generation traffic. A likely evolutionary scenario could be as follows. Assume that a network provider has a first generation SONET ring network. It could install new fiber, capable of

handling dense wavelength division multiplexing²⁷ (DWDM), along the same path as the first generation fiber. After the DWDM equipment is installed and tested, some number of λ s (equal in number and direction to the first generation discrete fibers) of the new fiber is separated and used to feed the first generation ADMs. Since these first generation ADMs cannot know that they are being served by λ s instead of discrete fibers, they continue to function in exactly the same manner as before, including protection features. The other λ s on the new fiber, however, can be used for second generation network techniques.

Perhaps because of this evolutionary requirement, many second generation networks are being configured similar to first generation networks, i.e., rings. Mesh networks are receiving more attention and we may see more of these as second generation networks reach a higher level of maturity.

The primary difference is that the add/drop multiplexer (ADM) becomes an optical add/drop multiplexer (OADM), as shown in Figure 21. Instead of dealing with SONET channels, the OADM deals with λ s, utilizing techniques such as fiber Bragg gratings to separate the λ s and micro-mirrors²⁸ (see Figure 22) to route them. Technically, the input to and output from an OADM is a wavelength of light. Any traffic within the λ must be dealt with outside the OADM.

Fiber Bragg Gratings

In 1978 it was discovered that optical fiber doped with germanium was photosensitive to ultraviolet (UV) light. The photosensitivity caused a change in the index of refraction, n . Photosensitivity was not exhibited at the longer wavelengths used for communications.

Building on this, scientists learned how to set up interference patterns in the UV light, thereby causing a periodic variation in n within the fiber. This pattern is known as a fiber Bragg grating (or filter) and can be used in place of bulk optics, allowing much smaller implementation of DWDM, with significantly less optical loss.

Fiber Bragg gratings have many applications in optical networks. See [Oth99] for more detail.

²⁷ It's not clear when wavelength division multiplexing becomes "dense" but it's probably when the λ s are spaced 200 GHz or less.

²⁸ Mechanical mirrors are unlikely to be the final solution for λ routing and switching – they're just too big and cause losses as the light leaves and re-enters the optical medium. Components based on guided wave optics are much more likely to be successful in commercial products.

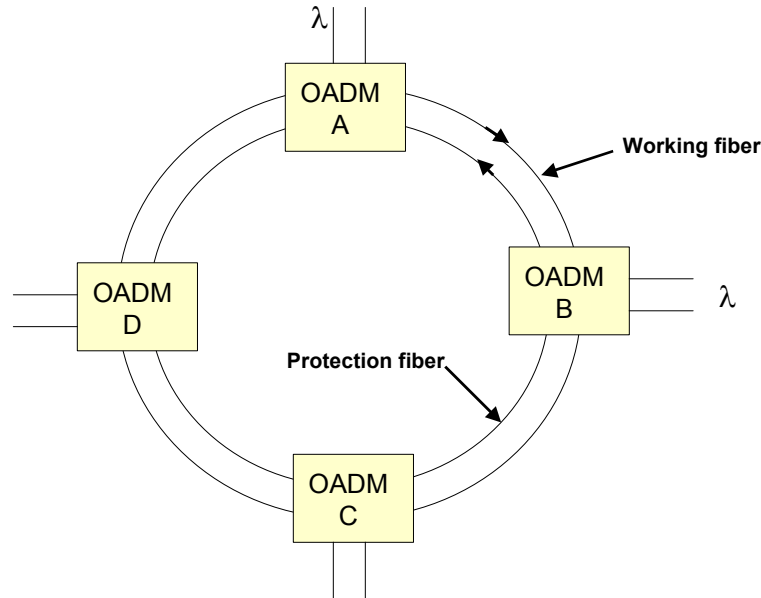


Figure 21: A second generation optical network, showing the use of optical add/drop multiplexers. The input and output of the network is λ s rather than digital data.

At one extreme, a completely photonic network could carry any traffic in a λ - analog or digital, and without any concern for the framing or timing of the digital data. Like all ideals, there are problems with implementing it in practice.




Figure 22: Example micro mirror used to redirect λ s in an all-optical network. Left, view of a single mirror; right, an array of mirrors. (source: Lucent LambdaRouter product)

The primary problem is that all optical networks are analog and subject to the cumulative impairments of analog systems, including optical nonlinearity, chromatic dispersion, amplifier induced four wave mixing, and others. At some point, the signal must be converted to electronic and retimed, reshaped and

retransmitted (3R regeneration – 2R is reshape and retransmit and can be done in the optical domain). An all-optical network will be limited by the capacity of the worse case route.

Additionally, an all-optical network must standardize technology choices at the outset and cannot easily take advantage of improvements in fiber or narrower optical spacing. For example, in the first generation network, new technology (e.g., DWDM with closer channel spacing) can be introduced on a link-by-link basis. In an all-optical network, certain new technology would be extremely difficult to introduce because it would have system wide effects.

Even with today's state-of-the-art optical technology, global-scale, or even national-scale, all optical networks are not feasible or practically attainable. Until optical components are developed which can achieve 3R regeneration, networks will have to utilize electronics to compensate for the analog impairments which accumulate in the analog domain.

Additionally, the network operator must be able to monitor the channel for errors in order to provide the specified quality of service. Today, all that can be done on a link is to monitor the optical power and perhaps the signal to noise ratio, which does not guarantee that digital data is being transmitted at an acceptable error rate. If the customer is experiencing an excessive error rate, it is very difficult in an all-optical network to determine where the errors are occurring. One way to deal with this problem is to allow electrical regeneration in the network (so called "*opaque networks*" [Bal95]) to isolate optical sections of multiple spans. This is equivalent to breaking the all-optical network into "islands" of optical switching, connected by electronic regenerators.

An approach to mitigate (but not solve) the problem of accumulated analog impairments and signal monitoring is to use forward error correction. There are two types of forward error correction proposed for use in optical networks, "in-band" and "digital wrappers." The in-band technique only works with SONET/SDH OC-48/STM-16 framed data. A BCH-3 code derived from a (8191, 8152) code is applied across the bits in a single row and the redundant bits (39 bits) are placed in unused locations in the overhead columns of the same or next row (note that there are no unused octets in rows 1 and 4). Eight codes are applied across each row, one for each bit of the octets (FEC1 is across all bit ones of all the octets in a row, FEC2 is across all bit twos of all octets in a row, etc.). Since each code can correct three bits, the interleaved eight codes can correct up to a 24 bit error burst. For higher line rates, the OC-48 streams are 16 bit interleaved. The proposed specifications for the in-band FEC are included in the revision of G.707 which was approved at the February 2001 ITU meeting in Geneva.

The digital wrapper is based on G.975 which defined an FEC structure for submarine cables. The terrestrial system is defined in the new recommendation G.709. The digital wrapper essentially takes the digital bit stream and breaks it into subframes, calculating a Reed Solomon (RS) code over the subframe. The code is a RS(255,239) code. One octet is taken from the 239 octet payload for framing, monitoring, and network management, providing about a 15/14 data rate. See Figure 23 which shows the format of one 255 octet subframe of the digital wrapper.




Figure 23: A digital wrapper subframe, utilizing a RS(255,239) code. One octet is taken from the 239 octets for framing and operations, administration, management and provisioning (OAM&P).

In the G.709 recommendation, 16 subframes are multiplexed by octet interleaving. The way this is done is shown in Figure 24. Note that the octets are transmitted from top to bottom, left to right, so the first octet is the framing/OAM&P octet. The next octet would consist of the first octet of the second subframe, all the way to the 16th subframe. After all the first octets are transmitted, the next octet will be the first payload octet of the first subframe, followed by the same octet in the next subframe, etc. This interleaves the data over the 16 sets of subframes to form a row.

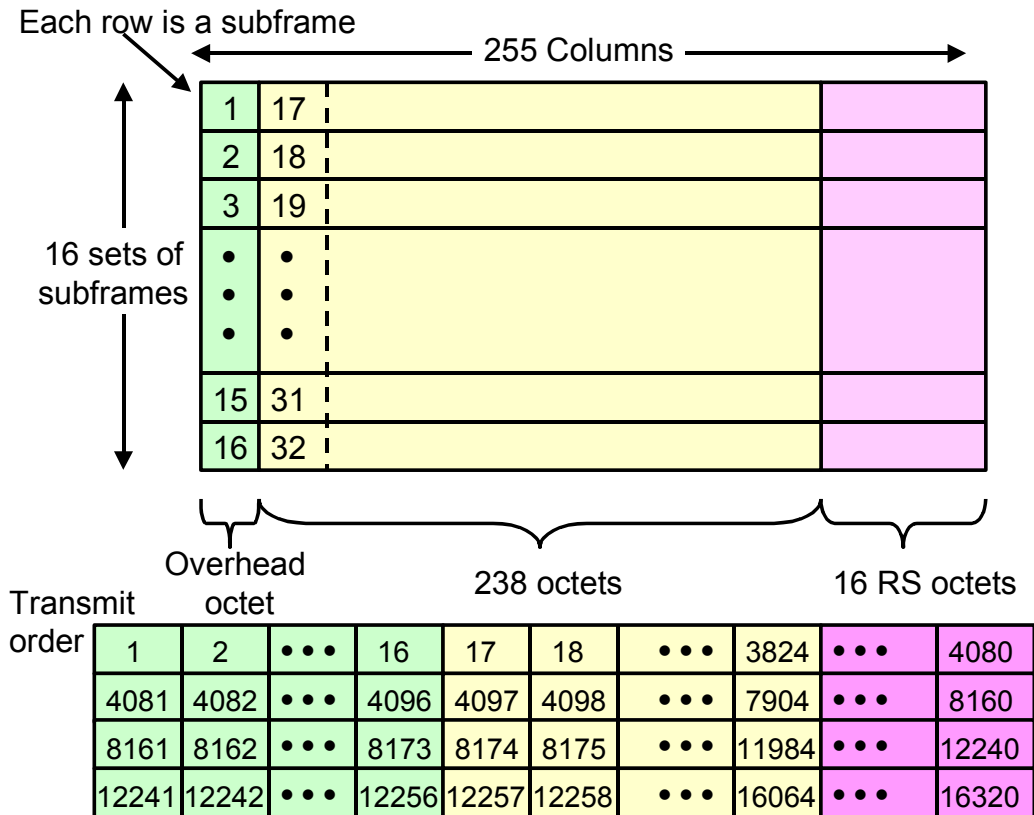


Figure 24: In the new recommendation G.709, 16 subframes are interleaved to form a row. Four rows then form a frame.

G.709 calls for a frame of 4 of these 16 interleaved rows. Figure 24 attempts to show this at the bottom of the figure by showing the four rows of the frame. Note that the subframes are interleaved but the rows are not – they are transmitted one after another in order to minimize delay and the amount of memory required by the encoder and decoder.

Framing characters are required and G.709 defines them to be the same as SONET/SDH (A1/A2, 0xf628). This requires that a scrambler be applied over the frame, except for the framing octets. The scrambler is defined in G.709.

The digital wrapper function is inserted between the data source/sink and the transceiver mux/demux function. See Figure 25 which shows a digital wrapper function placed between a SONET framer and the mux/demux function. Since this diagram implies SONET traffic, the equipment represented in the figure is set up to handle legacy traffic over digital wrapper.




Figure 25: Placement of the digital wrapper function in the communications chain. The digital wrapper and SONET framer are usually each a single chip, while the transceivers are usually separate Tx and Rx functions.

For OC-48, the use of the RS code increases the data rate from 2.488 Gbps to 2.666 Gbps. For OC-192, the data rate is increased from 9.953 Gbps to 10.709 Gbps. The increased data rate adversely affects the error rate, but the RS code provides gain. The net effect is to improve the BER of the line by perhaps 5 dB, which is a significant gain. Note that forward error correction, whether in-band or digital wrapper, does not solve all of the problems of an all-optical network – it simply mitigates the problems of analog impairments and provides some method of tagging and monitoring the signal along the route.

The frame-level digital wrapper looks a lot like a “SONET Lite” frame, prompting objections that the digital wrapper proponents are attempting to re-invent SONET. While there’s some merit to this argument, there are differences between the two. The digital wrapper technique is designed to be totally transparent to the traffic, but is obviously aimed at data traffic which has become the majority of network traffic. SONET was designed as a TDM technique, intended to carry voice channels. Digital wrapper frames are a fixed length, no matter what the line rate, while SONET frames are transmitted every 125 μ s, no matter the line rate. Also, the digital wrapper technique is an “overspeed” technique which maintains

the original SONET line rates (OC-3, OC-192, etc.), allowing support of legacy communications techniques.

At any point in the network, the operator can tap into the optical fiber (causing some loss of signal, perhaps 0.5 dB) and monitor the error rate²⁹. Alternately, the signal can be converted to electronic (at so-called OEO (optical-electronic-optical) points) to do 3R amplification and gain access to the monitoring and network management channel. The optical advocates see this being done only at the edge of the network, where connection is made to another network. If done only at the edges of the network, one λ must be dedicated to OAM&P and decoded at each OADM and other points where monitoring and control is required.

The introduction of a digital wrapper means that the second generation optical network cannot transport analog data, but this is not a real limitation. It also introduces certain limitations on the data rates within λ s. Since the digital wrapper logic must code and decode the data, it must be able to operate at the bit rate within the λ . This means that users will be restricted as to their use of a λ , although they could put any digital data desired into it – SONET, gigabit Ethernet, ATM over fiber, or any other protocol.

Routing within an all-optical network can become quite complex. For example, in Figure 21, if OADM A wishes to route traffic to OADM C (assume many more OADMs so that the alternate channel through OADM D is not appropriate), the network management system must know all of the λ s in use between OADM B and OADM C before making the channel (λ) assignment.

Also, as technology increases the number of λ s that a fiber can carry, and as the number of fibers served by an OADM goes up, the cost of being able to drop any λ at any OADM becomes prohibitively expensive. For example, assume a fiber is carrying 100 λ s. Any OADM will have, at a minimum, four fibers attached to it (one transmit and one receive, on each side), and could easily have many more than four. But even four fibers with 100 λ s each would require 400 sets of components to separate or combine the λ s, plus 200 sets of switching components (perhaps micromirrors). This is likely to be expensive and delicate. To deal with this problem, λ s are subdivided into bands, and only certain bands are provided with the ability to drop at a particular OADM. This further restricts the allocation of λ s for traffic.

All in all, there is much work to be done before practical all-optical networks are widespread.

Lambda Spacing

Commercial WDM systems are available which provide 80 to 100 λ s with 100 GHz spacing (these systems extend beyond the C-band to the L-band). The ITU specifies spacing with equal frequency increments, rather than equal wavelength increments. See G.692.

Researchers are also experimenting with 50 and 25 GHz spacing.

Note that the on-off keying (OOK) modulation of the laser light is equivalent to two level amplitude modulation. Since amplitude modulation requires a bandwidth equal to two times the modulation frequency, and since the modulation frequency is one half the bit rate, 10 Gbps operation requires a bandwidth of 10 GHz, establishing the lower limit of 10 GHz on λ spacing for OC-192 operation. OC-768 operation will require spacing of more than 40 GHz.

Spacing must be greater than the minimum because of tolerances in the laser frequency, frequency drift over the laser lifetime, and the requirements of real world filters.

²⁹ G.709 defines bit interleaved parity (BIP) octets for simple error checking.

Metropolitan (and Access) Optical Networks

The definition of “Metropolitan Optical Networks” varies across the industry. Many industry participants segment the overall network into backbone, metropolitan (meaning primarily local interoffice), and access. In this section, I address both the metropolitan and access segments, sometimes without making a clear distinction between the two. I do this because I believe that the two are interrelated and cannot be properly examined in isolation.

It’s difficult to estimate the market size of these segments but some industry observers see the metropolitan/access network segment as being ten times the size of the backbone segment (in dollars). When examining component costs, however, these same experts estimate that components for the metropolitan/access segment must be significantly less than the cost of backbone components. This means that the unit volumes in the metropolitan/access segments will be substantially greater than the unit volumes in the backbone segments. These unit volume and costs requirements will be the primary drivers for the technologies selected for these segments.

There’s little doubt that fiber will eventually be used as an access technology – no other technology offers the bandwidth and data rate that fiber does. As demand for bandwidth continues, fiber will have to be brought closer to the customer, whether it be a business or residential customer. However, it’s clear that business customers will be the leaders in the use of fiber for access. Businesses have the need for high data rates and connection to many other sites, including remote offices and other businesses (for B2B activities).

The concept of Metropolitan Optical Networks (MON) is just emerging and many techniques are being developed for it and for access. Certainly, any technique implemented must support legacy communications. Companies will have equipment and systems already installed. They will not convert overnight to any new communications technique, even if it is less expensive.

One conversion technique, of course, is to convert one service to the new architecture while leaving the remaining services on the existing access technologies. Over time, the legacy technologies can be converted to the fiber access. Even then, the fiber access must support legacy technologies such as T1, T3, ATM and SONET.

As more communications services are concentrated on the fiber access (even multiple fibers) protection becomes more important. Redundant electronics and multiple access fibers with separate routing will almost surely be used between the business and the network. The equipment at the corporate facility will function like an ADM or TM in this respect, monitoring the line and switching automatically when failures are detected.

It’s also clear that the access fiber must support WDM techniques. While WDM equipment is expensive, it is also extremely expensive to lay fiber in metropolitan areas. Once it’s laid, the service provider will do everything possible to get the maximum capacity out of the fiber. While the initial services will likely be provisioned with only a single wavelength, capacity will be added through WDM rather than by laying additional fibers for access to the business. This may be done by using λ s with wide spacing³⁰, allowing lower costs transmitter and receiver equipment. As additional capacity is needed, and as equipment costs decline, dense wavelength spacing can be utilized.

³⁰ Since the access link should be quite short, not requiring any amplification, both the 1310 nm and the 1550 nm bands could be utilized, allowing very wide spacing of λ s and lower cost λ separation equipment.

The access fiber could utilize the 1550 nm band for traffic that is to be optically routed over a second generation optical network, while the 1310 nm band could be utilized for traffic which is to be electrically terminated at the access point.

Since the access portion of the network is so large, it will be important to make it inexpensive. To as large a degree as possible, this portion of the network will be kept passive. The ideal situation would be for the access portion to consist simply of passive optical fibers and components, without EDFAs, with protection being provided by equipment in the corporation and in the feeder portion of the network (equipment in a central office or equivalent). There are ways to improve the performance of the optical fiber, including digital wrappers with Reed Solomon codes, so active devices in the access network should not be required.

The major architectural question for the future is how corporations will use bandwidth. One alternative is that corporations will continue to use private bandwidth, with essentially a transparent connection between the two end points. An example of this is a leased line between a corporation and a remote office. This architecture would encourage development of second generation optical networks where end to end links are provided and circuits can be set up as needed, perhaps even for short periods of time. This architecture is essentially a continuation of the circuit switched architecture of the traditional telephone network.

The other alternative is that corporations will use shared bandwidth, examples of which are frame relay, ATM and IP networks. Corporations are familiar with this model and have gained a level of comfort with it. Additionally, new techniques of encryption and authentication, such as is used with virtual private networks (VPNs), are becoming available which provide an extremely high level of protection.

Generally, this shared bandwidth model provides higher levels of efficiency than the circuit switched model, and should be expected to lead to lower cost access than the circuit switched model. Given these trends, my prediction is that corporations will move towards a single high speed "pipe" which provides access to a variety of end points, with switching at the higher protocol layers³¹.

No matter what model is used, there is the important question of what protocols will be used in the metropolitan network. SONET provides tools for management and for a natural transition to the backbone network. But even if SONET is used, we still need to determine the data protocol, since SONET does not frame data.

My belief is that Ethernet will migrate into the network, especially gigabit Ethernet and 10 gigabit Ethernet. This is a natural progression of what's happening inside the corporation today. Ethernet is distributed to every desktop from an Ethernet switch. Switches are interconnected using higher speed Ethernet. For Internet access, a router is connected to the backbone Ethernet and a serial connection (perhaps frame relay or ATM) is made to the Internet.

Suppose that the corporation wished to interconnect remote offices and also provide connections to other corporations in the area which it did business with. It could lease fiber "lines," using them to build an Ethernet network in the metropolitan area. Using today's technology, it could run 100Base-FX (100 Mbps Ethernet over fiber) or 1000Base-LX (1000 Mbps Ethernet over fiber). It might decide for logistical reasons to locate the network Ethernet switch at a location separate from the corporate headquarters, but central to the leased circuits. Using this all-Ethernet network, the corporation could communicate at high speed with its remote offices and business partners, without having to do additional protocol conversions.

³¹ Circuit switching will not disappear, but the growth will be in the shared bandwidth area.

Now, suppose that someone came to the corporation and offered to supply the same service at lower costs by making that Ethernet network available to a large group of companies, thereby spreading the costs. They might even offer 10,000 Mbps Ethernet (10 gigabit Ethernet) access. Ethernet would now be part of the “public” network.

There are many advantages to this scenario. Ethernet is simple and therefore cheap. Much of the routing could be done at layer 2, which is extremely fast. Costs would be low since the equipment would be very similar to the equipment used inside the corporation, leading to higher volume. And just as in the corporation, routing could also be done at layers above layer 2, providing QoS.

This traffic would, of course, still be IP traffic – the IP frames would be carried in the Ethernet frames. It’s just that routing and switching would be done on whichever was most convenient, the Ethernet MAC address³² or the IP address.

This concept would allow Ethernet to migrate into the network. Initially, there would be small “islands” of Ethernet in local areas, and all traffic beyond those islands would be passed to an Internet router, exactly as is done inside the corporation today. Over time, the Ethernet domains could grow, until it was possible to communicate via Ethernet with sites around the globe.

This is not the only concept being pursued for metropolitan networks. Several startups are pursuing alternate technologies but only time will tell which, if any, are successful

Summary

*True luminescence
Is seen on a summer’s eve
In a firefly’s light.
Ruth E. Riter
Photonics Haiku*

This paper attempted to introduce some of the technologies used in optical networks. By necessity, it only described the major concepts – there’s a lot more for those who wish to really understand the various technologies utilized in modern optical networks. Additionally, the paper attempted to identify some emerging technologies in this expanding market.

As Dickens’ observed in the nineteenth century, “[These are] the best of times, [these are] the worst of times.” We are faced with extremely rapid change in the communications network, offering both threat and opportunity. As an industry, we must move rapidly to exploit the opportunities available to us.

³² Note that Ethernet MAC addresses are globally unique. There will never be a duplicate unless done maliciously.

Appendix A – Modulation in Optical Communications

After the original publication of this paper, several people asked for more detail on the format of the modulation of the light in optical systems.

The basic technique of modulation is on-off keying (OOK) where the light is turned off and on to signal the presence of binary ones and zeros. But there's different ways of doing this.

The simplest way is to provide light for the duration of a “one” bit time and turn the light off for the duration of a “zero” bit time. This technique is known as non-return to zero (NRZ). Optical NRZ is different than electrical NRZ because there's no negative light as there is negative voltage. The diagram indicated as NRZ in Figure 26 would be known as return to zero (RZ) in the electrical world.




Figure 26: Two common methods of modulating the light signal on optical fibers: non-return to zero (NRZ) and return to zero (RZ).

The other technique is to turn the light on for a one bit, but to turn it off before the bit time, allowing the signal to go dark. The reason for using RZ modulation is to assure transitions for long strings of one bits. With NRZ modulation if a long string of ones was to be transmitted, the light would be turned on and would stay on for the duration of the one bits. Depending on the length of the string, this could cause the receiver to lose timing and be unable to decode the bits correctly.

While RZ solves this problem for long strings of ones, it still permits long strings of zeros and reduces the number of photons seen by the photodetector by half for each one bit. This is a serious problem so RZ is not used for long distance transmission to any great degree.

Another problem is DC balance. DC balance is another term taken from electrical communications but has a slightly different meaning in optical communications. In optical communications DC balance means that the average transmitted power is constant. This is important at the receiver because it affects the decision threshold.

For long distance communications NRZ is typically used along with some technique to minimize long strings of ones and zeros and to maintain DC balance.

The two techniques typically used to avoid long strings of ones and zeros are line codes and scrambling. I won't describe either of these techniques in detail but will note that scrambling does not add any additional bits to the transmitted stream but can not guaranteed to eliminate long strings of ones or zeros. However, it can make the probability of their occurrence low. It also cannot guarantee absolute DC balance.

Line codes map some number of bits, say 8 bits, to a subset of a larger number of bits, perhaps 10 bits. Since I have more numbers with 10 bits than I can have with 8 bits, I can select a subset of the 10 bit numbers and assign one to each 8 bit number. If I do this properly, I can guarantee a sufficient number of transitions (I can limit the number of sequential ones or zeros) and I can guarantee DC balance. The price I pay is that I must transmit more bits; in this case 10 bits must be transmitted for each 8 information bits.

Bibliography

- [Bal95] Bala, K., R.R. Cordell, E. L. Goldstein, “The case for opaque multiwavelength optical networks,” *Proceedings of the IEEE/LEOS Summer Topical Meeting on Global Information Infrastructure*, Keystone, CO, Aug 95.
- [Dos99] Doshi, B. T., S. Draviad, E. J. Henrandez-Valencia, W. A. Matragi, M. A. Qureshi, J. Anderson, and J. S. Manchester, “A simple data link protocol for high-speed packet networks,” *Bell Labs Technical Journal*, January-March 1999, pp 85-104.
<http://www.lucent.com/minds/techjournal/jan-mar1999/pdf/paper05.pdf>
- [Bre99] Breen, Jim. “Overview of Fibre Optics.” html page, 1999?
<http://www.csse.monash.edu.au/~jwb/subjects/cse4891/fo/fo.html>
- [Gla00] Glass, Alastair M., David J. DiGiovanni, Thomas A. Strasser, Andrew J. Stentz, Richard E. Slusher, Alice E. White, A. Refik Kortan, and Benhamin J. Eggleton. “Advances in Fiber Optics,” *Bell Labs Technical Journal*, January-March 2000, pp 168-187.
<http://www.lucent.com/minds/techjournal/jan-mar2000/pdf/paper10.pdf>
(Very good paper on the present state and future directions in fiber optics.)
- [Gor97] Goralski, Walter J. *SONET, A Guide to Synchronous Optical Networks*. McGraw- Hill, 1997. (Basic information, sort of a “SONET for Dummies.”)
- [Has95] Hasegawa, Akira and Yuji Kodama. *Solitons in Optical Communications*. Clarendon, 1995. (Advanced text. Dense and difficult to read.)
- [Hec99] Hecht, Jeff. *City of Light*. Oxford University Press, 1999. (Good historical information on the development of optical communications)
- [Hen01] Henderson, P. Michael. “*Fundamentals of SONET/SDH*”. 2001. (Available at <http://members.cox.net/michael.henderson/>)
- [Hen01b] Henderson, P. Michael. “*Forward Error Correction in Optical Networks*.” 2001. (Available at <http://members.cox.net/michael.henderson/>)
- [Her99] Hernandez-Valencia, Enrique. “A Simple Data Link (SDL) framing protocol for high-speed optical packet networks.” Optical Internetworking Forum (OIF) contribution OIF99.043.0 (available at <http://www.oiforum.com> but you need an ID and password).
- [Mca99] McAdams, Larry and Iain Verigin. “A proposal to use POS for the OIF physical layer up to OC-192c.” Optical Internetworking Forum (OIF) contribution OIF99.002.2 (available at <http://www.oiforum.com> but you need an ID and password).
- [Oth99] Othonos, Andreas, and Kyriacos Kalli. *Fiber Bragg Gratings: Fundamentals and Applications in Telecommunications and Sensing*. Artech, 1999. (An excellent introductory book on fiber Bragg gratings – recommended.)
- [Ram98] Ramaswami, Rajiv, and Kumar N. Sivarajan. *Optical Networks*. Morgan Kaufman, 1998. (Excellent text on all aspects of Optical Networks – highly recommended.)

- [Sei98] Seifert, Rich. *Gigabit Ethernet*. Addison-Wesley, 1998. (Good book on Ethernet in general, in addition to gigabit Ethernet.)
- [Sil96] Siller, Curtis A., Jr. and Mansoor Shafi. *SONET/SDH*. IEEE Press, 1996. (Broad coverage but best read by those who already have some knowledge of SONET/SDH.)
- [Sim19] Simpson, W. RFC 1619 – PPP over SONET. Internet Engineering Task Force, 1994. <http://ftp.isi.edu/in-notes/rfc1619.txt>
- [Sim61] Simpson, W. RFC 1661 – The Point-to-Point Protocol. Internet Engineering Task Force, 1994. <http://ftp.isi.edu/in-notes/rfc1661.txt>
- [Sim62] Simpson, W. RFC 1662 – PPP in HDLC-like Framing. Internet Engineering Task Force, 1994. <http://ftp.isi.edu/in-notes/rfc1662.txt>
- [Str00] Strand, John. “*Fundamental limits of optical transparency.*” Powerpoint presentation sent with personal correspondence in July 2000. Presentation dated 1998.
- [Sun99] Sun, Yan, Atul K. Srivastava, Jianhui Zhou, and James W. Sulhoff. “Optical Fiber Amplifiers for WDM Optical Networks.” Bell Labs Technical Journal, January-March 1999, pp 187-206. Available at <http://www.lucent.com/minds/techjournal/jan-mar1999/pdf/paper10.pdf>