Optical Communication Networks

EE-654

Lecture-1

Spring 2016

Propagation of Signals in Optical Fiber

Optical fiber is a remarkable communication medium compared to other media such as copper or free space. An optical fiber provides low-loss transmission over an enormous frequency range of at least 25 THz—even higher with special fibers—which is orders of magnitude more than the bandwidth available in copper cables or any other transmission medium. For example, this bandwidth is sufficient to transmit hundreds of millions of phone calls simultaneously, carry about a million high definition TV (HDTV) video streams, The low-loss property allows signals to be transmitted over long distances at high speeds before they need to be amplified or regenerated. It is because of these two properties of low loss and high bandwidth that optical fiber communication systems are so widely used today.

Still, the fiber itself does impose physical limitations that must be taken into account in network design. The goal of this chapter is to provide an understanding of the three phenomena that determine fiber transmission limits: loss, nonlinear effects, and *dispersion*. Dispersion is the phenomenon whereby different components of a signal travel at different velocities. In most cases, dispersion limits the data rate of a digital signal by spreading signal pulses over time. In Chapter 5 the interaction of loss, nonlinearity, and dispersion in designing advanced systems will be discussed.

1.1 Loss and Bandwidth Windows

The loss incurred by propagating down a fiber can be modeled easily as follows: the output power P_{out} at the end of a fiber of length L is related to the input power P_{in} by

$$P_{out} = P_{in} e^{-\alpha L}.$$

Here the parameter α represents the fiber attenuation. It is customary to express the loss in units of dB/km; thus a loss of α_{dB} dB/km means that the ratio P_{out}/P_{in} for $L = 1$ km satisfies

$$10 \log_{10} \frac{P_{out}}{P_{in}} = -\alpha_{dB}$$

or

$$\alpha_{dB} = (10 \log_{10} e)\alpha \approx 4.343\alpha.$$

The two main loss mechanisms in an optical fiber are *material absorption* and *Rayleigh scattering*. Material absorption includes absorption by silica as well as the impurities in the fiber. The material absorption of pure silica is negligible in the entire 0.8–1.6 μm band that is used for optical communication systems. The reduction of the loss due to material absorption by the impurities in silica has been very important in making optical fiber the remarkable communication medium that it is today. The loss has now been reduced to negligible levels at the wavelengths of interest for optical communication—so much so that the loss due to Rayleigh scattering is the dominant component in today's fibers in all three wavelength bands used for optical communication: 0.8 μm, 1.3 μm, and 1.55 μm. Figure 2.1 shows the attenuation loss in silica as a function of wavelength. We see that the loss has local minima at these three wavelength bands with typical losses of 2.5, 0.4, and 0.25 dB/km. (In a typical optical communication system, a signal can undergo a loss of about 20–30 dB before it needs to be amplified or regenerated. At 0.25 dB/km, this corresponds to a distance of 80–120 km.) The attenuation peaks separating these bands are primarily due to absorption by the residual water vapor in the silica fiber.

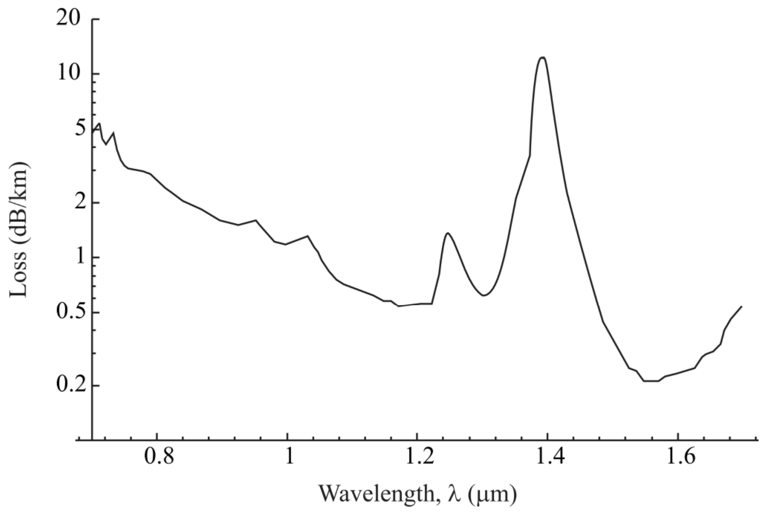

Figure 2.1 Attenuation loss in silica as a function of wavelength. (After [Agr97].)

The bandwidth can be measured in terms of either wavelength $\Delta\lambda$ or frequency Δf. These are related by the equation

$$\Delta f \approx \frac{c}{\lambda^2}\Delta\lambda.$$

This equation can be derived by differentiating the relation $f = c/\lambda$ with respect to λ. Consider the long wavelength 1.3 and 1.5 μm bands, which are the primary bands used today for optical communication. The usable bandwidth of optical fiber in these bands, which we can take as the bandwidth over which the loss in decibels per kilometer is within a factor of 2 of its minimum, is approximately 80 nm at 1.3 μm and 180 nm at 1.55 μm. In terms of optical frequency, these bandwidths correspond to about 35,000 GHz! This is an enormous amount of bandwidth indeed, considering that the bit rate needed for most user applications today is no more than a few tens of megabits per second.

The usable bandwidth of fiber in most of today's long-distance networks is limited by the bandwidth of the erbium-doped fiber amplifiers (see Section 3.4) that are widely deployed, rather than by the bandwidth of the silica fiber. Based on the availability of amplifiers, the low-loss band at 1.55 μm is divided into three regions, as shown in Figure 2.2. The middle band from 1530 to 1565 nm is the conventional or C-band where WDM systems have operated using conventional erbium-doped fiber amplifiers. The band from 1565 to 1625 nm, which consists of wavelengths longer than those in the C-band, is called the L-band and is today being used in high-capacity WDM systems, with the development of gain-shifted erbium-doped

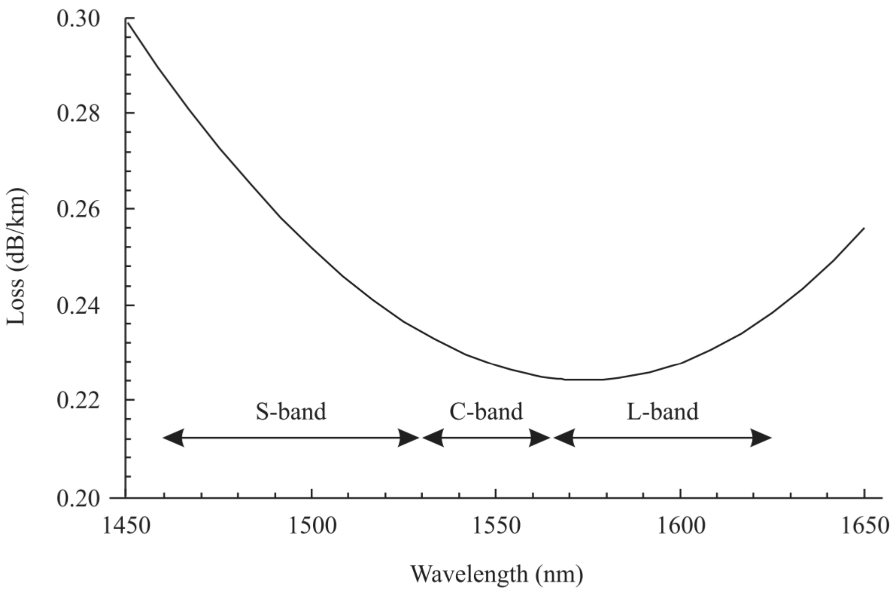

Figure 2.2 The three bands, S-band, C-band, and L-band, based on amplifier availability, within the low-loss region around 1.55 μm in silica fiber. (After [Kan99].)

amplifiers (see Section 3.4) that provide amplification in this band. The band below 1530 nm, consisting of wavelengths shorter than those in the C-band, is called the S-band. Fiber Raman amplifiers (Section 3.4.4) provide amplification in this band.

Lucent introduced the AllWave single-mode optical fiber, which virtually eliminates the absorption peaks due to water vapor. This fiber has an even larger bandwidth, and is useful where there are no erbium-doped fiber amplifiers.

As we saw earlier in this section, the dominant loss mechanism in optical fiber is Rayleigh scattering. This mechanism arises because of fluctuations in the density of the medium (silica) at the microscopic level. We refer to [BW99] for a detailed description of the scattering mechanism. The loss due to Rayleigh scattering is a fundamental one and decreases with increasing wavelength. The loss coefficient α_R due to Rayleigh scattering at a wavelength λ can be written as $\alpha_R = A/\lambda^4$, where A is called the Rayleigh scattering coefficient. Note that the Rayleigh scattering loss decreases rapidly with increasing wavelength due to the λ^{-4} dependence. Glasses with substantially lower Rayleigh attenuation coefficients at 1.55 μm are not known. In order to reduce the fiber loss below the current best value of about 0.2 dB/km, one possibility is to operate at higher wavelengths, so as to reduce the loss due to Rayleigh scattering. However, at such higher wavelengths, the material absorption of silica is quite significant. It may be possible to use other materials such as fluorozirconate ($ZiFr_4$) in order to realize the low loss that is potentially possible by operating at these wavelengths [KK97, p. 69].

1.2 Bending Loss

Optical fibers need to be bent for various reasons both when deployed in the field and particularly within equipment. Bending leads to “leakage” of power out of the fiber core into the cladding, resulting in additional loss. A bend is characterized by the *bend radius*—the radius of curvature of the bend (radius of the circle whose arc approximates the bend). The “tighter” the bend, the smaller the bend radius and the larger the loss. The bend radius must be of the order of a few centimeters in order to keep the bending loss low. Also, the bending loss at 1550 nm is higher than at 1310 nm. The ITU-T standards specify that the additional loss at 1550 nm due to bending must be in the range 0.5–1 dB, depending on the fiber type, for 100 turns of fiber wound with a radius of 37.5 mm. Thus a bend with a radius of 4 cm results in a bending loss of < 0.01 dB. However, the loss increases rapidly as the bend radius is reduced, so that care must be taken to avoid sharp bends, especially within equipment.

1.3 Intermodal Dispersion

An optical fiber consists of a cylindrical *core* surrounded by a *cladding*. The cross section of an optical fiber is shown in Figure 2.3. Both the core and the cladding are made primarily of silica (SiO_2), which has a refractive index of approximately 1.45. The *refractive index* of a material is the ratio of the speed of light in a vacuum to the speed of light in that material. During the manufacturing of the fiber, certain impurities (or dopants) are introduced in the core and/or the cladding so that the refractive index is slightly higher in the core than in the cladding. Materials such as germanium and phosphorus increase the refractive index of silica and are used as dopants for the core, whereas materials such as boron and fluorine that decrease the refractive index of silica are used as dopants for the cladding. As we will see, the resulting higher refractive index of the core enables light to be guided by the core, and thus propagate through the fiber.

1.4 Multimode and Single-mode Fiber

Just as there are different grades of copper cables, there are many grades of optical fiber. The most fundamental divide is between *single-mode* and *multimode* fiber. The difference between the two is so profound it is often better to think of them

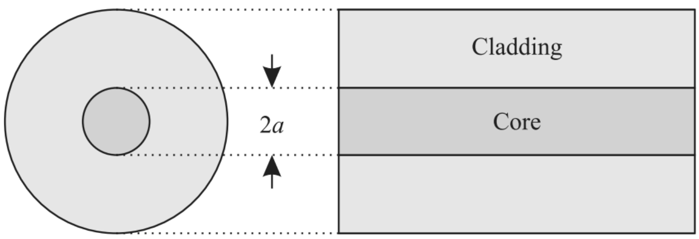

Figure 2.3 Cross section and longitudinal section of an optical fiber showing the core and cladding regions. *a* denotes the radius of the fiber core.

as completely different media types, almost as different as copper and fiber. The majority of this book is concerned with single-mode fiber because that is the medium for networks of any length above a few hundred meters. However, multimode fiber will be discussed in this section.

Typical multimode fiber has a core much larger than a wavelength of light. As a result, a simple geometric optics view can be used to describe its overall behavior, which we present in Section 2.2.1. Multimode fiber carries hundreds of modes, which can be thought of as independently propagating paths of the optical signal. Signals on different modes have different velocities. This creates *intermodal dispersion*. In most situations, dispersion leads to broadening of signal pulses, which correspond to data bits. In a communication system, this leads to the overlap of pulses representing adjacent bits, distorting the signal. This phenomenon is called *Inter-Symbol Interference* (ISI).

Single-mode fiber has a core on the same scale as a wavelength that restricts itself to a single "fundamental" spatial core. This eliminates intermodal dispersion. Hence, single-mode fiber is used for the highest bandwidth and longest distance transmission. However, since its core is on the same scale as a wavelength, a true electromagnetic wave treatment as presented in Section 2.3.1 is necessary to understand its behavior.

1.5 Geometrical Optics Approach

We can obtain a simplified understanding of light propagation in optical fiber using the so-called *ray theory* or *geometrical optics* approach. This approach is valid when the fiber that is used has a core radius a that is much larger than the operating wavelength λ. These are *multimode* fibers, and first-generation optical communication links were built using such fibers with a in the range of 25–100 μm and λ around 0.85 μm.

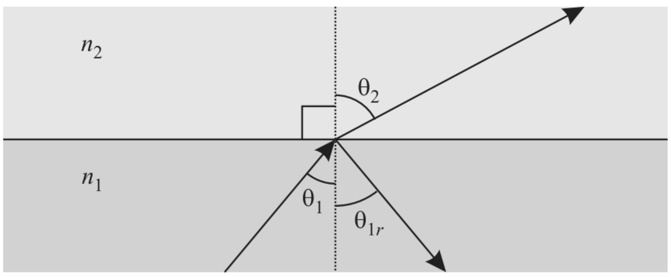

Figure 2.4 Reflection and refraction of light rays at the interface between two media.

In the geometrical optics approach, light can be thought of as consisting of a number of "rays" propagating in straight lines within a material (or medium) and getting reflected and/or refracted at the interfaces between two materials. Figure 2.4 shows the interface between two media of refractive index n_1 and n_2. A light ray from medium 1 is incident on the interface of medium 1 with medium 2. The *angle of incidence* is the angle between the *incident ray* and the normal to the interface between the two media and is denoted by θ_1. Part of the energy is reflected into medium 1 as a *reflected ray,* and the remainder (neglecting absorption) passes into medium 2 as a *refracted ray.* The *angle of reflection* θ_{1r} is the angle between the reflected ray and the normal to the interface; similarly, the *angle of refraction* θ_2 is the angle between the refracted ray and the normal.

The laws of geometrical optics state that

$$\theta_{1r} = \theta_1$$

and

$$n_1 \sin\theta_1 = n_2 \sin\theta_2. \tag{2.1}$$

Equation (2.1) is known as *Snell's law*.

As the angle of incidence θ_1 increases, the angle of refraction θ_2 also increases. If $n_1 > n_2$, there comes a point when $\theta_2 = \pi/2$ radians. This happens when $\theta_1 = \sin^{-1} n_2/n_1$. For larger values of θ_1, there is no refracted ray, and all the energy from the incident ray is reflected. This phenomenon is called *total internal reflection*. The smallest angle of incidence for which we get total internal reflection is called the *critical angle* and equals $\sin^{-1} n_2/n_1$.

Simply stated, from the geometrical optics viewpoint, light propagates in optical fiber due to a series of total internal reflections that occur at the core-cladding interface. This is depicted in Figure 2.5. In this figure, the coupling of light from the medium outside (taken to be air with refractive index n_0) into the fiber is also shown.

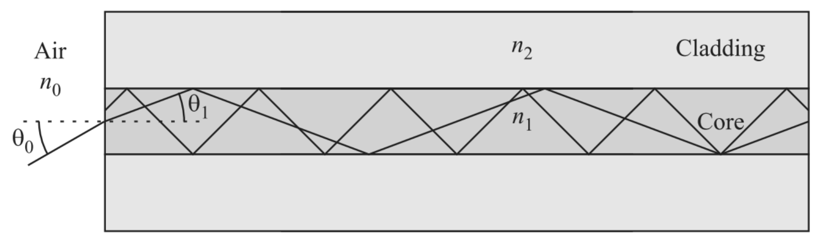

Figure 2.5 Propagation of light rays in optical fiber by total internal reflection.

It can be shown using Snell's law (see Problem 2.1) that only those light rays that are incident at an angle

$$\theta_0 < \theta_0^{\max} = \sin^{-1} \frac{\sqrt{n_1^2 - n_2^2}}{n_0} \tag{2.2}$$

at the air-core interface will undergo total internal reflection at the core-cladding interface and will thus propagate. Such rays are called *guided rays,* and $\theta_0^{\max}$ is called the *acceptance angle.* The refractive index difference $n_1 - n_2$ is usually small, and it is convenient to denote the fractional refractive index difference $(n_1 - n_2)/n_1$ by Δ. For small Δ, $\theta_0^{\max} \approx \sin^{-1} \frac{n_1\sqrt{2\Delta}}{n_0}$. As an example, if $\Delta = 0.01$, which is a typical value for (multimode) fiber, and $n_1 = 1.5$, a typical value for silica, assuming we are coupling from air, so that $n_0 = 1$, we obtain $\theta_0^{\max} \approx 12°$.

1.6 Bit Rate-Distance Limitation

Owing to the different lengths of the paths taken by different guided rays, the energy in a narrow (in time) pulse at the input of the fiber will be spread out over a larger time interval at the output of the fiber. A measure of this time spread, which is called *intermodal dispersion*, is obtained by taking the difference in time, δT, between the fastest and the slowest guided rays. Later we will see that by suitably designing the fiber, intermodal dispersion can be significantly reduced (graded-index fiber) and even eliminated (single-mode fiber).

We now derive an approximate measure of the time spread due to intermodal dispersion. Consider a fiber of length L. The fastest guided ray is the one that travels along the center of the core and takes a time $T_f = Ln_1/c$ to traverse the fiber, c being the speed of light in a vacuum. The slowest guided ray is incident at the critical angle on the core-cladding interface, and it can be shown that it takes a time $T_s = Ln_1^2/cn_2$ to propagate through the fiber. Thus

$$\delta T = T_s - T_f = \frac{L}{c}\frac{n_1^2}{n_2}\Delta.$$

How large can δT be before it begins to matter? That depends on the bit rate used. A rough measure of the delay variation δT that can be tolerated at a bit rate of B b/s is half the bit period $1/2B$ s. Thus intermodal dispersion sets the following limit:

$$\delta T = \frac{L}{c}\frac{n_1^2}{n_2}\Delta < \frac{1}{2B}. \tag{2.3}$$

The capacity of an optical communication system is frequently measured in terms of the *bit rate–distance product*. If a system is capable of transmitting x Mb/s over a distance of y km, it is said to have a bit rate–distance product of xy (Mb/s)-km. The reason for doing this is that usually the same system is capable of transmitting x' Mb/s over y' km providing $x'y' < xy$; thus only the product of the bit rate and the distance is constrained. (This is true for simple systems that are limited by loss and/or intermodal dispersion, but is no longer true for systems that are limited by chromatic dispersion and nonlinear effects in the fiber.) From (2.3), the intermodal dispersion constrains the bit rate–distance product of an optical communication link to

$$BL < \frac{1}{2}\frac{n_2}{n_1^2}\frac{c}{\Delta}.$$

For example, if $\Delta = 0.01$ and $n_1 = 1.5(\approx n_2)$, we get $BL < 10$ (Mb/s)-km. This limit is plotted in Figure 2.6.

Note that θ_0^{max} increases with increasing Δ, which causes the limit on the bit rate–distance product to decrease with increasing Δ. The value of Δ is typically chosen to be less than 1% so as to minimize the effects of intermodal dispersion, and since θ_0^{max} is consequently small, lenses or other suitable mechanisms are used to couple light into the fiber.

1.7 Controlling Intermodal Dispersion: Graded-index Multimode Fiber

Thus far, we have assumed that the fiber is a *step-index* fiber since the variation of the refractive index along the fiber cross section can be represented as a function with a step at the core-cladding interface. In practice, however, multimode fibers have more

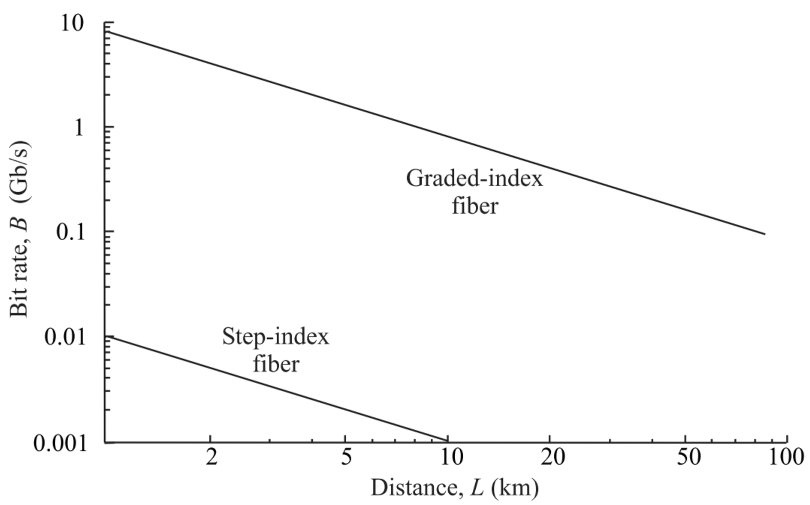

Figure 2.6 Limit on the bit rate–distance product due to intermodal dispersion in a step-index and a graded-index fiber. In both cases, $\Delta = 0.01$ and $n_1 = 1.5$.

sophisticated *graded-index* profiles designed to reduce the intermodal dispersion. The refractive index decreases gradually, or continuously, from its maximum value at the center of the core to the value in the cladding at the core-cladding interface.

This has the effect of reducing δT because the rays traversing the shortest path through the center of the core encounter the highest refractive index and travel slowest, whereas rays traversing longer paths encounter regions of lower refractive index and travel faster. For the optimum graded-index profile (which is very nearly a quadratic decrease of the refractive index in the core from its maximum value at the center to its value in the cladding), it can be shown that δT, the time difference between the fastest and slowest rays to travel a length L of the fiber, is given by

$$\delta T = \frac{L}{c}\frac{n_1 \Delta^2}{8}.$$

Assuming that the condition $\delta T < 1/2B$, where B is the bit rate, must be satisfied, we get the following limit on the bit rate–distance product of a communication system employing graded-index fiber:

$$BL < \frac{4c}{n_1 \Delta^2}.$$

For example, if $\Delta = 0.01$ and $n_1 = 1.5$, we get $BL < 8$ (Gb/s)-km. This limit is also plotted in Figure 2.6 along with the limit for step-index fiber.

Table 2.1 A comparison of multimode fiber. Effective modal bandwidth (EMB) and overfilled launch bandwidth (OFLBW) correspond to laser and LED sources, respectively.

Fiber Type	Also Known As	Core Diameter (Microns)	EMB–Distance at 850 nm (MHz-km)	OFLBW–Distance at 850/1300 nm (MHz-km)
OM1	FDDI Grade	62.5	*Not Applicable*	200/500
OM2		50.0	*Not Applicable*	500/500
OM3	Laser Optimized	50.0	2000	1500/500
OM4	Laser Optimized	50.0	4700	3500/500

1.8 Solitons

Solitons are narrow pulses with high peak powers and special shapes. The most commonly used soliton pulses are called *fundamental solitons*. The shape of these pulses is shown in Figure 2.25.

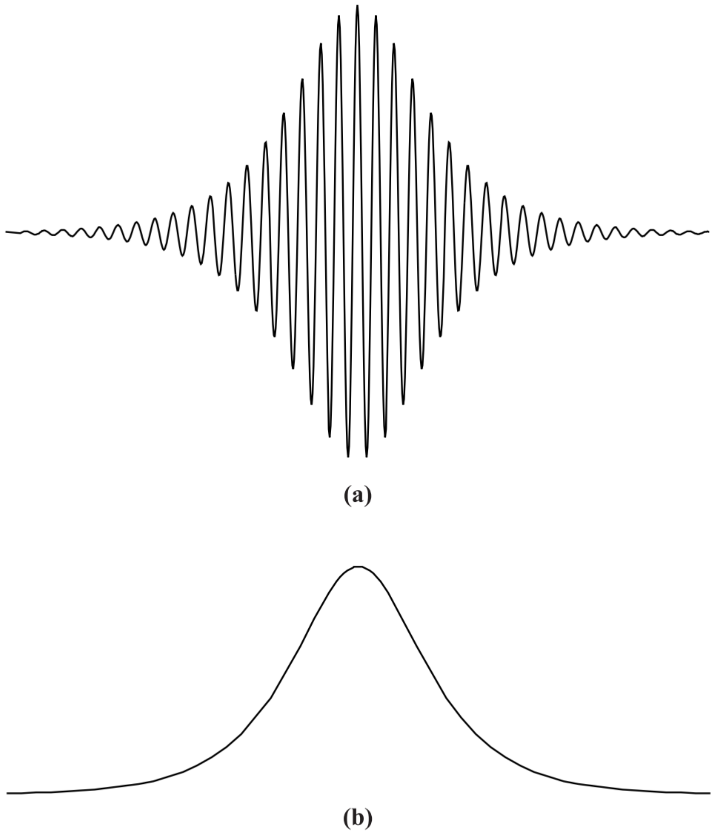

Figure 2.25 (a) A fundamental soliton pulse and (b) its envelope.

The significance of solitons for optical communication is that they overcome the detrimental effects of chromatic dispersion completely. Optical amplifiers can be used at periodic intervals along the fiber so that the attenuation undergone by the pulses is not significant, and the higher powers and the consequent soliton properties of the pulses are maintained. Solitons and optical amplifiers, when used together, offer the promise of very high-bit-rate, repeaterless data transmission over very large distances.

1.9 Optical Fiber Technologies

We will discuss two fiber types that are not traditional glass fibers. The first is designed by having periodic structures, and the second uses plastic material.

Photonic Crystal Fiber

In previous sections we have seen how dopants and fiber profile can be engineered to reduce loss, dispersion and nonlinearity, for better transmission. There is another category of fiber designs that is not limited by bulk material properties. As in semiconductors, engineers can create sometimes startling properties that do not exist in bulk materials by playing with periodic structures and defects in periodic structures within the fiber. These fiber designs are called *photonic crystal fibers* (PCFs).

PCFs were first demonsrated in 1996 and have been an active area of research since then. Some of the properties that can be created are dispersion, nonlinearity, and even negative refractive index (e.g., according to Snell's law, as illustrated in Figure 2.4, if the refractive index is positive, the rays are refracted on the opposite side of the normal on entering the material, but negative refractive index means rays will be refracted on the same side).

PCF enables a number of functions in fiber, some of which are relevant to dispersion compensation, amplification, and wavelength conversion by nonlinear optics. The PCF structures for fiber have been in two dimensions. We should note that the associated science and fabrication of PCF has extended beyond fiber to materials for other devices and that structures in three dimensions are being explored as well.

All the fiber types described below are "holey" fibers, in which the glass material is laced with a carefully designed pattern of holes. Figure 2.26 shows what crosssections of holey fiber may look like. We can see that structures have a pattern in two dimensions. A common way to make such a fiber is to bundle together tubes of glass and then to draw out the fiber. Fabrication of these holes and maintaining the precision of the design while the fiber is drawn continues to be a challenge that limits their use to specialty purposes.

However, there are two very different classes of holey fiber, which work on different physical principles: *index guiding* and *photonic bandgap*. (Hybrid versions that take advantage of both effects also exist.)

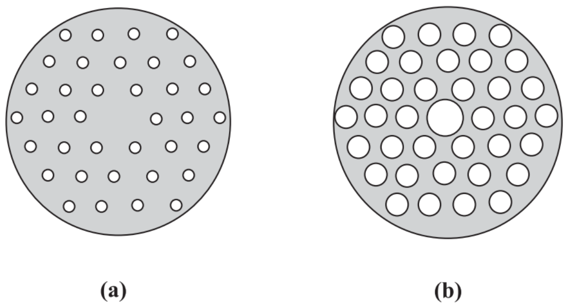

Figure 2.26 Two examples of the crosssection of holey fibers.

Index Guiding

The fibers we will discuss next have a periodic structure but do not rely on the periodicity to provide fiber guiding. They still use index guiding. To get a flavor of what these structures allow, suppose we would like to reduce the fiber's bending loss (see Section 2.1.1). One way to reduce loss is to confine the light more strongly to the core by increasing the index difference between the core and cladding. However, (2.12) implies that if the core size is kept constant, then the fiber will become multimoded at a longer cutoff wavelength, possibly even at the operational wavelength of interest. Conversely, if the core size is decreased to maintain the desired cutoff wavelength, the mode size would shrink to the point that it becomes impractical to effectively connect the fiber to other components, for example, by splicing or other connector technologies.

Now suppose air holes are introduced into the cladding of the fiber. If these holes are small enough that the optical mode "sees" only the average index, the effective cladding index—an average between the original cladding material index and the air index of unity—is dramatically lowered. The result is a very large index difference.

In reality, the index "seen" by the mode is more complex than simply the average of air and glass indices since there is no way to completely remove the effect of the periodic structure. Also note that the periodic structure can be optimized by a designer so that its effects will improve the fiber's performance further. One of these effects is discussed in the next section on photonic bandgap fibers.

A holey fiber that has found commercial application is the Corning ClearCurve, which uses holes in a ring within the cladding. These holes are so small (several hundred nanometers in diameter) that the material is described as a "nanostructure." The ClearCurve fiber can be bent tightly (5 mm radius) with minimal loss. Cable made with the fiber can be handled like electrical wire when installed at residences for fiber-to-the-home, which is not possible with ordinary single-mode fiber.

Holey fiber can also have its holes filled with materials rather than air. For example, materials with high nonlinearity, including gases and liquid crystals, have been introduced into the holes. These hybrids allow the designer to combine desired properties of fiber guiding with a host of other material properties.

Photonic Bandgap

An air-guided fiber has a periodic array of holes running longitudinally down the fiber, and these holes define a guiding structure. The "core" is defined by a defect or extra hole such that the guided mode exists mostly in air. Figure 2.26(b) shows an example with a hole in its center.

Note that since the "core" is mostly in air, it should have a lower index than the surrounding cladding. Thus, this fiber does not exploit total internal reflection to confine the light to the core. Instead, it uses the periodic structure of the holes. The structure creates a *photonic bandgap*, which is a range (or band) of wavelengths for which propagation is forbidden. The principle behind it is the same as that used in Bragg gratings, which is covered in Section 3.3.3. A Bragg grating is a periodic perturbation in the propagation medium, usually a periodic variation of the index of refraction. For the fiber, the periodicity of the hole structure in the cladding destructively interferes with light of certain wavelengths that attempts to penetrate it. The periodicity in the cladding is designed to have a bandgap for the range of wavelengths used in operation. Then wavelengths within the range are confined to the core. This phenomenon is called a photonic bandgap because it is analagous to an electronic bandgap found in semiconductors.

In practice, the tolerances required have kept air-guiding fiber from commercial use to date. The demonstrated losses have been higher than conventional transmission fiber, and the manufacturing difficulty is considerably greater.

Plastic Optical Fiber

In today's home networks, many types of media are being used to connect increasingly high-speed data feeds between set-top boxes, computers, storage, and various pieces of audiovisual equipment. These include various forms of copper (coax installed by cable operators, existing phone lines, existing power lines), wireless, and even fiber.

One advantage plastic optical fiber has over glass fiber for home networking is simplicity of termination, which can be done with a penknife or plastic-melting tools. Another material-based difference from glass fiber is its long-term reliability under sustained bends. Plastic optical fiber is able to creep (i.e., gradually deform over very long times) to relieve strain and hence does not suffer the chemical surface changes that afflict strained glass.

Plastic optical fiber has been in the home for decades. For example, the Sony/Philips Digital Interconnect Format (S/PDIF) interfaces are used to carry audio signals between devices and stereo components. The physical medium can be optical fiber using the TOSLINK® (TOShiba-Link) standard. This application uses step-index polymethyl methacrylate (PMMA) fibers whose 1 mm total diameter consists of a 980 micron diameter plus 10 micron thick cladding ring. The core index is 1.49 and the cladding index is 1.42. The bit rate-distance product is 10 MHz-km. Because of the short distances of the applications, the bandwidth limitations have typically not come from the fiber but from the speed of the transmitters, which are 650 nm LEDs. The material does not transmit in the infrared, thus disallowing the use of 850 nm VCSEL transmitters used in gigabit per second data communications.

Perfluorinated graded-index fiber (POF) is designed to reduce the material absorption loss at 850 nm wavelength, so that the fiber can be used with VCSELs for high-speed home networking. The highest bit rate-distance products are obtainable with smaller cores. As the core gets smaller, the design becomes similar to that of silica-based multimode fiber. The trade-off is that increasing the bit rate-distance product reduces both the mechanical tolerances for connectors and bend insensitivity.

Summary

Understanding light propagation in optical fiber is key to appreciating not only the significant advantages of using optical fiber as a propagation medium but also the problems that we must tackle in designing high-bit-rate WDM systems. We started by understanding how light propagates in multimode fibers using a simple ray theory approach. This introduced the concept of pulse broadening due to multimode dispersion and motivated the use of single-mode fibers. After describing the elements of light propagation in single-mode fibers, we studied the limitations imposed on optical communication systems due to the pulse-broadening effects of chromatic dispersion.

Although dispersion is the most important phenomenon limiting the performance of systems at bit rates of 2.5 Gb/s and below, nonlinear effects become important at higher bit rates. The main nonlinear effects that impair high-speed WDM transmission are self-phase modulation and four-wave mixing. We studied the origin of these, as well as other nonlinear effects, and briefly outlined the constraints on optical communication systems imposed by them. We will return to the system limitations of both dispersion and nonlinearities when we discuss the design of optical transmission systems in Chapter 5.

We also studied the new types of fibers that have been introduced to mitigate the effects of dispersion and nonlinearities. We then discussed solitons, which are special pulses designed to play off dispersion and nonlinearities against each other to achieve high-bit-rate, ultra-long-haul transmission.

We also discussed new types of multimode fiber, and novel fiber types such as holey and plastic fibers.