Chapter 5
Fiber Optics

Optical Fibers

As fiber optic cables fall in price, they are being used more and more for communications. In a nutshell, electrical signals are converted into a light beam (possibly visible, but more likely infrared) with an LED (light-emitting diode) or laser, the light beam is sent through a fiber made of glass or plastic, and then converted back into an electrical signal with a photodiode or phototransistor. Thus the fiber cable simply replaces the copper wire.

The optical fiber, however, has a number of features that make it very much better for some applications. For one, its attenuation is much, much lower than that of cable, especially at high bandwidths. It is entirely practical to send high-speed data or analog signals some 30 or more miles without needing to amplify or regenerate signals along the way.

Optical fibers also offer immunity to electrical interference. They do not leak energy outward, are extremely difficult to tap, and do not accept interference from other nearby signals. They are also immune to nearby lightning strikes, although this is sometimes compromised by cables that contain both optical fibers as well as electrical conductors inside the same cable.

Fiber optic cables can be made of either glass or plastic. A typical cable generally looks like Fig. 5-1, and consists of three parts:

a. The *core* is at the center. Made of very pure and clear glass or plastic, the core passes the light beam from one end of the fiber to the other.

b. The *cladding* is a second layer of glass or plastic, wrapped around the core. This layer is made of a slightly different glass or plastic, one that lets light travel just a bit faster than the core.

c. The *jacket* covers the outside of the fiber and protects it. In some fibers, additional jacket layers, including perhaps even metal sheathing, may provide additional protection.

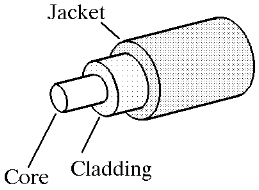

Fig. 5-1. Optical fiber construction

To understand how this works, imagine that we are laying out an automobile race course. In order to keep cars from straying off course, we construct it as shown in Fig. 5-2. In the middle is the race course itself, constructed of a very smooth surface designed for average traction. This is where we want each car to drive.

On each side of the racing lane we put a second lane made of a slightly different material, one which provides slightly better traction. This is labelled the speedup lane in the figure. Finally, at the very edge, just in case a car gets too far off course, we'll put a layer of sand to stop it so it can't wander back into the course and get in someone else's way.

So now look at car 1, which has somehow gotten slightly off course. As soon as it enters the speedup lane, the left wheels will have slightly better traction than the right wheels, so its left side will wind up travelling just a bit faster. This is going to turn it slightly to the right, so it goes back onto the racing lane. Car 2, on the other hand, is so far off course that it too will turn a bit to its right, but not enough to return to the course. It will eventually get stuck in the sand, and that will keep it out of trouble.

Thus the speedup lane keeps the cars from straying off course, without any intervention on their part. You could almost put a blindfolded driver into a car, and the course itself would keep him on track.

The optical fiber works exactly the same way. The cladding is made from a glass that lets light travel just a bit faster than the core.

DETOUR

In terms of physics, the core has a slightly higher *index of refraction* than the cladding; this makes light travel slightly slower in the core than in the cladding. The equation that describes how the light bends at the boundary between the core and the cladding is called *Snell's Law*; it is really quite simple, but not necessary for us right now. The result, though is that when a ray of light tries to leave the core and enter the cladding, it is bent back just enough to re-enter the core. If it's like car 1, it returns back to the core and keeps going. But if it's

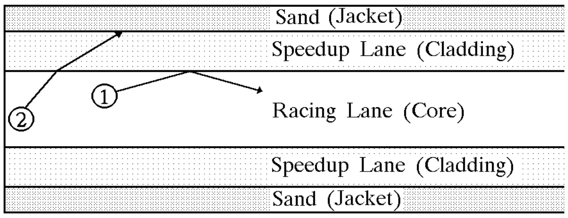

Fig. 5-2. Keeping a ray on track

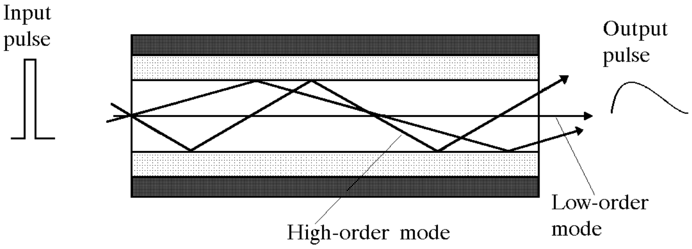

Fig. 5-3. Multi-mode fiber

like car 2, then it goes right through the cladding until it hits the jacket, and gets stopped.

It is important to realize that the process of “bouncing” off the cladding back into the core is not the same as reflection off a mirror. Reflection in a mirror always loses a small amount, since no mirror is perfect. *Refraction*, as it is called, is 100% efficient — there is no light lost in the process. The only light that is lost is due to any impurities in the glass.

END OF DETOUR

Fiber optic cables can actually be made from either glass or plastic. As such, they come in three types:

- All glass is best, and most expensive
- All plastic is worst, but cheapest
- Plastic-clad silica (PCS) is a compromise, where the core is quartz-like form of glass called silica, but the cladding is plastic.

The difference is primarily due to the fact that glass is clearer, and therefore significantly more light gets through the fiber. Because of the loss of light, plastic fibers are only usable in short lengths of a few hundred feet, whereas glass fibers can be dozens of miles long.

Transmission Modes

Referring to Fig. 5-2, we can see that a car which goes straight down the core will make better time than a car which constantly bounces from edge to edge. This is an extremely important concept in optical fibers as well, because it determines the dispersion of light.

Fig. 5-3 shows a *stepped-index* or *multi-mode* fiber. It is called stepped-index because there is a sudden change (mathematicians call it a step) in the index between the core and the cladding; it is called multi-mode because it permits many modes. The *mode* describes the path that a light beam takes: a *low-order mode* is one that goes directly through, taking as few bounces as possible. A *high-order mode* is one that takes many bounces before arriving at the far end. Obviously, a high-order mode light ray must travel a greater distance to the far end, and so arrives later than a low-order light ray. If the light is pulsed on and off, as shown at the left, the pulse arrives at the far end spread out, as shown at the right. The term *dispersion* describes how spread out the pulse is at the output.

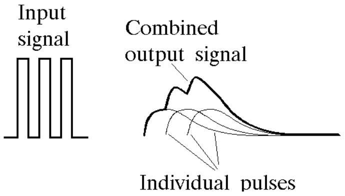

Fig. 5-4. Dispersion spreads pulses

The amount of dispersion depends on the length of the fiber. When a fiber is short, the difference between a low-order path and a high-order path is fairly small, so the total dispersion is fairly small as well. The longer the fiber, the greater the difference in path length, and therefore the greater the dispersion.

Fig. 5-4 shows why dispersion is a problem. Most fiber optic cables are used for digital signals; the bandwidth of such a cable depends on how many pulses can be sent per second. When a fast series of such pulses is sent through the cable, dispersion spreads out each pulse as shown. The pulses then overlap, making it harder to properly separate and detect them. When dispersion makes adjacent pulses blend together too much, the only solution is to separate the pulses farther apart. This limits the number of pulses that can be sent per second. In other words, the bandwidth of the fiber depends on its dispersion, which in turn depends on its length — the greater the length, the lower the bandwidth.

A *graded-index* fiber is a partial solution to the problem. Rather than having a sudden change in the index of refraction between the core and the cladding, the graded-index fiber has a gradual change in the index. This has two effects. First, it results in the light rays gradually bending back toward the middle, rather than abruptly “bouncing” back and forth. This tends to

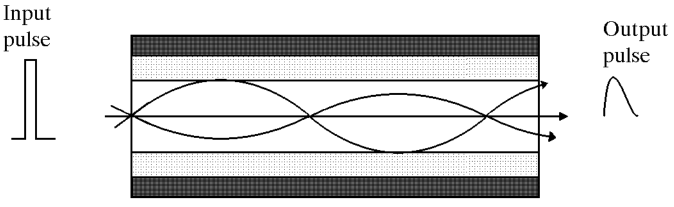

Fig. 5-5. Graded-index fiber

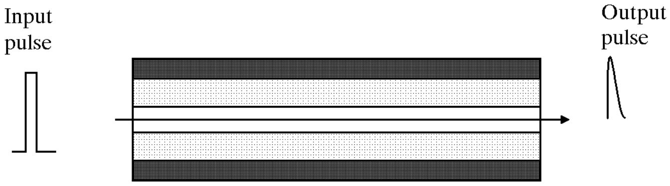

Fig. 5-6. Single-mode fiber

gradually straighten out the light path, causing more of the rays to travel along low-order modes rather than high-order paths. A more important effect is caused by the fact that the index of refraction is highest in the middle of the fiber, and gradually decreases as you go outward. Light rays travelling toward the edges of the fiber therefore move faster than those in the middle. Even though they travel a longer distance, the faster speed lets them catch up with the lower mode rays in the center. The result is less dispersion.

The best solution to the dispersion problem is with a *single-mode* fiber; a fiber which has only a single possible mode of light transmission. Just as two out-of-phase electrical signals can cancel, so two out-of-phase light rays can cancel. Because of this interference effect, there are only certain specific light paths (modes) in a fiber which are possible; all others cancel out because of interference. If the core is made thin enough, the only mode that is possible is the one straight-through mode in the middle. The resulting single-mode fiber, shown in Fig. 5-6, thus has much less dispersion and therefore a higher bandwidth.

You can understand, though, that a very thin single-mode fiber will let through much less light than a thicker multi-mode or graded-index fiber. For that reason, single-mode fiber generally requires a glass core which, because of its clarity, attenuates the light less. Plastic fibers, on the other hand, are usually thicker to let through more light. This increases dispersion and thus reduces their bandwidth, but this is not much of a problem because their attenuation is so high that they can only be used in shorter lengths; at these lengths, dispersion is not as bad as it would be in real long lengths.

Fiber characteristics

Let us next look at some characteristics of actual fibers of various kinds, listed in Table 5-1.

Core and cladding diameter

Fiber diameters are given in μm or micrometers — 10^{-6} meters. To put this into context, note that a typical human hair is about 50 μm thick, so the thinnest fibers are about 1⁄10 of the thickness of a hair. Fiber dimensions are usually stated as two numbers separated by a slash; that is, a 50/125 fiber has a 50 μm core diameter and a 125 μm cladding diameter.

In general, the plastic fibers are much thicker than the glass fibers, although multi-mode glass fibers can also be quite thick. The thinner fibers obviously must be much clearer to let through enough light.

Attenuation

The attenuation determines how long a cable can be before the signal becomes too small to be useful. It is rated in dB of signal loss per kilometer of length. The best glass fibers have losses below 0.5 dB per kilometer, whereas some of the plastic fibers have losses of as much as 400 dB per km.

The best glass fibers are amazingly clear — a 6 km length of glass single-mode fiber, about 4 miles, would have a loss of about 3 dB. To realize just how amazing this is, think of a piece of window glass 4 miles thick. With ordinary glass, you probably would not even be able to detect light through it. With a fiber, a 3 dB loss means that half of the light entering it comes out the far side!

Even with 30-mile lengths (about 50 km) of a good glass fiber, the loss is only some 20 or so dB — a loss of

Table 5-1. Optical fiber characteristics

Type	Core Dia (μm)	Cladding dia (μm)	Attenuation (dB/km)	BW MHz·km	NA
Single mode glass	5	125	0.5	1000+	small
	8	125	0.4	1000+	small
Graded index glass	50	125	4	400	.20
	63	125	7	250	.29
	85	125	6	200	.26
	100	140	5	20	.30
Multi-mode glass	200	380	6	25	.27
	300	440	6	20	.27
Multi-mode PCS	200	550	10	15	.30
	800	900	6	20	.30
Multi-mode Plastic	980	1000	400	20	.50

99% of the light power, but still within a reasonable range. Plastic fibers, on the other hand, can have huge losses. At 400 dB per km, 20 dB loss would occur in about $\frac{1}{20}$ km, or about 150 feet. Because of this high loss, the maximum usable length of a plastic fiber is a few hundred feet — long enough to go within a small building, but not much more.

It is interesting to compare fiber attenuation with the attenuation in copper cables. The attenuation of copper cables increases with the frequency of the signal. For example, good quality RG-8 with a foam dielectric has about 1.2 db loss per 100 feet at 50 MHz, about 2 db loss at 150 MHz, and 3 db loss at about 300 MHz.

Let's look only at the loss at 50 MHz, since that is close to the 47.736 MHz data rate of a T3 line, a common data speed in optical fibers. Since there are about 3300 feet in a kilometer, a 1 km length of RG-8 would have about 40 db loss at 50 MHz. Even one of the best and most expensive coax cables ($\frac{7}{8}$-inch diameter hardline) would have a loss of 12 db per km at 50 MHz. Compare this with the 0.5 db per kilometer loss of a good glass optical fiber, and you can clearly see the advantages of fibers over long distances.

Attenuation vs. Frequency

As we just mentioned, the attenuation in copper cables depends greatly on the frequency, and therefore affects the bandwidth. As the frequency increases, the attenuation goes up very fast, and this greatly limits the bandwidth.

When talking about fiber cables, however, we must be careful, because there are two frequencies to consider — the frequency of the light (which actually describes the light's color), and the frequency of the data (which describes how fast the light is turned on and off in the case of digital data, or how fast its intensity changes in the case of analog data.) Rather than use frequency, however, physicists use the wavelength since it is easier to measure (as you remember, the wavelength is the velocity divided by the frequency, so it is easy to convert back and forth.) The wavelength is measured in nanometers or nm (one nm is 10^{-9} meters), and the colors measure approximately as follows:

Ultraviolet: shorter than 400 nm
Violet: 400–450 nm
Blue: 450–500 nm
Green: 500–570 nm
Yellow: 570–590 nm
Orange: 590–610 nm
Red: 610–750 nm
Infrared: from about 750 to about 1400 nm

(Ultraviolet and infrared wavelengths are also considered to be light, but they are invisible to the eye.) Light's wavelengths are so short that they are greatly affected by molecular and atomic characteristics of materials. For example, Fig. 5-7 shows the attenuation of glass at various wavelengths. We see that the attenuation varies quite a bit with wavelength. At very short wavelengths (toward the left of the graph), light transmission is affected by individual atoms — the further left you go, the greater the attenuation, so these wavelengths are not very useful for optical fibers. At longer wavelengths (toward the right), light transmission is affected by certain combinations of atoms, and so there are specific wavelengths that are highly attenuated, while other nearby wavelengths get through with relative ease. As long as we stay away from these attenuation peaks, the red and infrared wavelengths will work best.

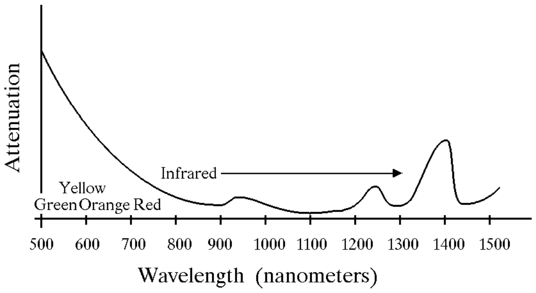

Fig. 5-7. Attenuation of glass fiber

This is fortunate, because LEDs and diode lasers are easier to build, more efficient, and cheaper at infrared and deep red colors. Glass fibers are generally used with infrared light between 790 and 1300 nm because that is where the attenuation is least. Plastic fibers, on the other hand, are usually used with red light at about 650 nm.

Note that it is possible to send several different colors through a fiber cable at the same time, as long as we choose the colors wisely. Thus it is possible to send two or more different signals through the cable. Still, this is not too common in longer cables, because feeding two or more color sources into the same fiber, and then separating them again with colored filters, introduces losses that limit the distance.

The attenuation of various colors is thus a separate topic from the concept of signal bandwidth.

Signal bandwidth

The bandwidth of a fiber is related to how fast we can turn a beam on and off, or more likely, how fast we can change its brightness (because even in digital applications, the beam is usually switched between bright and dim, rather than between bright and off.)

This speed is determined by two factors: the speed at which the light sources and detectors can work, and the dispersion in the fiber itself.

As we showed in Fig. 5-4, dispersion makes adjacent pulses blend together; they therefore have to be separated further apart. Since the dispersion increases with the length of a fiber, the pulse rate for a given fiber is inversely proportional to the distance. If the length (in km) is large, the frequency at which we can switch the beam (in MHz) is small, and vice versa. The product of these two quantities is therefore fairly constant for a given fiber, and is given in "MHz times km" or MHz·km. This number is listed in the bandwidth column of Table 5-1.

For example, the bandwidth for single-mode glass fiber is shown as 1000+ MHz·km (the + means that it is 1000 or greater.) Assuming a value of an even 1000, a 1-km length of fiber would have a bandwidth of 1000 MHz; a 10-km length would work up to 100 MHz, and so on.

Compared with 1000+, the rating of 20 MHz·kmfor multi-mode plastic fiber looks awful ... until we put it into context. Remember that, because of its attenuation, plastic fiber is limited to a small fraction of a kilometer. For instance, consider a plastic fiber $\frac{1}{20}$ km long; at that length, the bandwidth is still a respectable 400 MHz (since 400 MHz times $\frac{1}{20}$ km is 20 MHz·km.) Thus plastic fiber can also be used for high bandwidths, but clearly only for short distances. This helps to explain why no one bothers to make single-mode plastic fibers — there is simply no way to take advantage of the lower dispersion that this would provide.

NA or Numerical Aperture

The last column in Table 5-1 shows the *numerical aperture* for various fibers. The meaning of this measurement is shown in Fig. 5-8.

The dark line at the left in the figure shows a light ray which comes into the fiber at such a steep angle that it hits the cladding and stays there; it does not "bounce" off it back into the core. Clearly such a light ray is not going to make it to the far end. Furthermore, any other light ray that is at that angle or more will not make it through either.

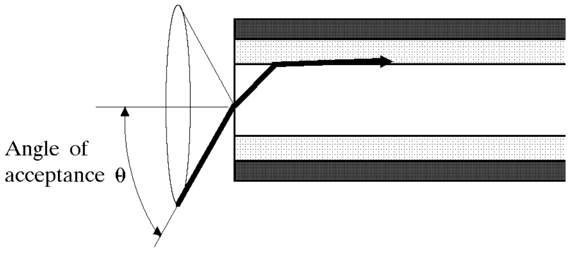

Fig. 5-8. Numerical Aperture - cone of acceptance

The result is that there is what we call an *angle of acceptance* θ, which defines the cone shown in Fig. 5-8. Light coming in within that cone will be transmitted through the cable; light coming outside that cone will not. The light coming out the far end of a fiber will generally come out in a cone of the same angle, though in some cases the higher-order light modes can be attenuated in a longer fiber, with the result that the light comes out in a narrower cone.

The numerical aperture NA is simply the sine of the angle θ, namely

$$\text{numerical aperture NA} = \sin \theta$$

From Table 5-1, we see that the NA is between 0.2 and 0.3 for most of the fibers shown, corresponding to angles θ of 12 to 18 degrees. Multi-mode plastic fiber's NA of 0.5 corresponds to 30 degrees, while single-mode glass fiber has acceptance angles of less than 10 degrees.

The NA is important when matching a fiber to the light source or light detector, or when connecting two fibers together. For example, a plain light bulb would be a very bad light source for a fiber, since its light comes from a large area, most of which would be outside the angle of acceptance.

Connectors and splices

Fiber cables can be connected together in various ways, but the connection must be done carefully to avoid light loss from one cable to the next. Splices and fiber optic connectors have to do two things: hold the two fibers firmly together end to end, and align the two cables so as much light as possible goes from one into the other.

Connections can be broken down into three types:

- Pluggable connectors
- Splices where the two fibers remain separate
- Splices where the two fibers are joined into one

There are a number of factors that come into play:

Distance between fibers. If the two fibers are separated from each other, light will escape the joint. Thus those splices where the two fibers actually join into one are best. But when the fibers remain separate, it is usually best that they do not touch, because actual physical contact between them can cause scratches which deflect or even reflect the light. Holding the two fibers firmly, keeping them aligned, yet not letting them touch obviously requires care. In some connectors, a liquid with a similar index of refraction is placed into the gap between the two fibers to provide a better light match.

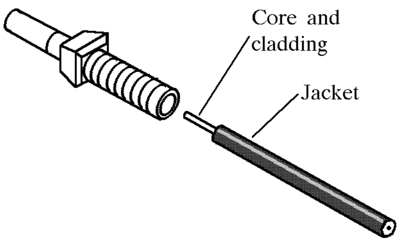

Fig. 5-9. One type of plug for plastic fiber

Side-to-side alignment. The two fibers must be aligned side-to-side to within a small fraction of the core diameter. For example, with a 5 µm core diameter, a side-to-side alignment error of more than one µm would be critical; with a 980 µm core, one µm would be irrelevant.

Core and cladding diameter. Clearly, connecting two identical cables is best. Their diameters can be different, however, as long as the smaller fiber feeds the larger one. In this case, the light from the smaller fiber enters only the middle of the larger one, but none is lost. If, on the other hand, a larger fiber feeds a smaller one, then some of the light coming out of the larger fiber "spills" outside the core of the second fiber, and light is lost.

Numerical Aperture. Again, connecting two identical cables is best, because the cones of acceptance then match. If the NA is different, however, then it is best that the fiber with the smaller NA feed the one with the larger NA; in this case, the light coming out of the first fiber comes out in a narrow cone, which fits into the wider cone of the second fiber. If the opposite is true — the larger NA feeds the smaller NA fiber — then some of the light coming out of the first fiber again "spills" outside the cone of acceptance of the second fiber, and light is lost.

Pluggable connectors

Pluggable connectors come in many types; Fig. 5-9 shows one kind of connector for plastic cable. Once the outer jacket is stripped from the inner core and cladding for about ¼ inch, the fiber is slipped into the connector. Metal teeth inside the connector keep the jacket from pulling back out. The connector's manufacturer provides a metal tool which slips over the narrow end of the connector, and guides a hot knife which is used to cut the excess fiber so it protrudes just slightly from the end. The flat side of the knife is then used to melt the end of the fiber flush with the end of the connector.

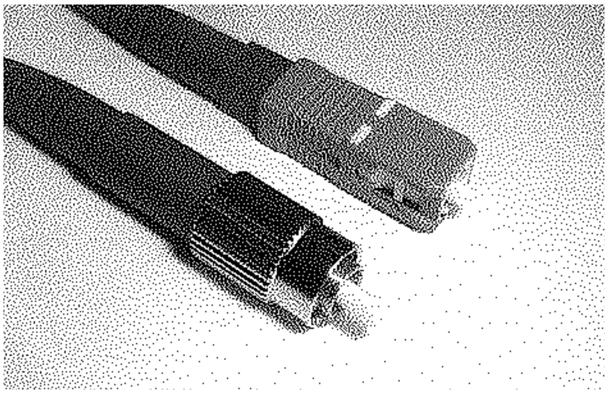

Fig 5-10. Single-mode FC and SC connectors

From this description, you can imagine that this is not a very accurate or delicate process. In general, plastic cables have so much loss that most cables are short; a few dB more or less is then not considered significant, and there is no need for super precision.

The situation is quite different with glass fibers, where it is usually more important to minimize losses. Whereas plastic fibers are just cut with a hot knife, glass fibers are carefully cleaved. That is, a narrow groove is scored into the glass, and then the glass is bent slightly until it snaps. The end of the fiber is then carefully polished, rather than just being melted with a hot knife. Sometimes, the end is polished flat; other times it has a slight convex curve. There are a number of connectors for single-mode glass fibers; two are shown in Fig. 5-10.

Splices where the two fibers remain separate

Splicing two fibers so that they remain separate (as opposed to being fused together), is also done in various ways. For example, a plastic splice very similar to Fig. 5-9 is available from the same manufacturer. It looks like a double-ended connector, with one fiber being slipped in from each end. As before, metal teeth hold the outer jacket to prevent the fiber from slipping back out. Since the connector itself is black plastic, it is impossible to see how close the two fibers are to each other. Fortunately, the multi-mode fibers are thick enough to tolerate slight misalignments.

Glass fibers, and especially single-mode fibers, require much more careful alignment. These splices are usually done with the aid of a special metal or plastic block. A small groove cut into the block holds the two fibers; once you position the two fibers in place, a cover is attached to the top to squeeze and hold the two fibers in place.

Fig. 5-11. A semi-automatic fusion splicer

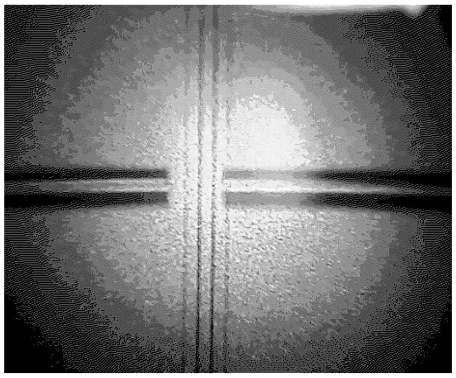

Fig. 5-12. Microscope display on semi-automatic splicer

Fig. 5-13. Fully automatic splicer

Fig. 5-14. TV display from the automatic splicer

Glass fibers can also be mechanically spliced by being held together in a small tube and glued with an optical glue.

Splices where the two fibers join together

Glass fibers are most often spliced with a *fusion splice*. Here the two fibers are carefully aligned together, and then an electric arc melts their ends and fuses them together into one. Finally, heat-shrinkable tubing is slipped over the splice to hold the fibers and make the assembly rigid.

In the early days of fibers, this was a purely manual operation done by carefully trained technicians. Splicing today, however, is so frequent that there are several commercial fusion splicers on the market which make the job almost trivial.

Fig. 5-11 shows one such unit made by Fujikura. This splicer projects an image of the fiber onto a small microscope screen; the technician then uses the controls to carefully align the fibers to get them into position, and then pushes a button to start the fusing process. Fig. 5-12 shows the display during the splicing process. This particular unit is fairly small and popular with users who have to make splices in uncomfortable surroundings, such as on top of a telephone pole.

Fig. 5-13 shows another fusion splicer. This one is somewhat larger and heavier, but it is fully automated. Once you cut and strip the fiber, you simply insert it into the machine, and it automatically aligns and fuses the fibers. Instead of a microscope screen, this unit has a built-in TV camera which shows the process on a small LCD display screen, complete with messages at the bottom to indicate progress. It can also display the image on an external TV monitor; Fig. 5-14 shows that image during splicing.

A well-done fusion splice can have as little as 0.2 dB loss; any fusion splice with a loss of 1 dB or more will generally be redone. A good non-fusion splice, on the other hand, seldom approaches even the 1 dB level.

Light Sources

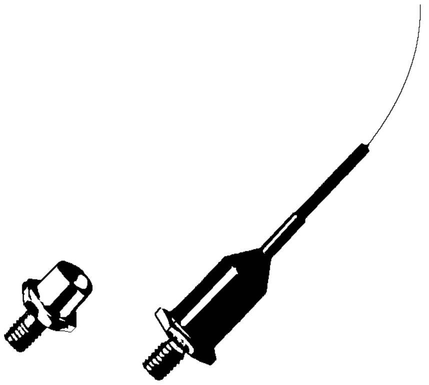

Fig. 5-15. Semiconductor laser diodes

The light source for a fiber optic system can be either a light-emitting diode (LED) or a semiconductor laser.

LEDs are perfectly adequate for short cables. They are much cheaper, simpler to use, and they also last longer.

Semiconductor laser diodes, however, have a number of advantages that outweigh their higher cost and complexity. Not only are laser diodes faster and brighter, but their main advantage, especially for longer runs of fiber, is that their light output is much more pure.

We have earlier discussed the problem of dispersion, and how dispersion limits the bandwidth of a fiber. What we did not say at that time was that the speed of light through a fiber depends slightly on the color of the light; that is, on its wavelength. The light from an LED, although it looks like a pure color (and is, in fact, much purer than the light from a colored light bulb), actually still consists of a range of wavelengths. The various wavelengths in a pulse of LED light will therefore arrive at the far end at different times, causing dispersion of the pulse. A laser, on the other hand, generates a very pure light with a very narrow range of wavelengths. Its light is therefore dispersed much less.

This is not important for lower bandwidths or short lines, but it is critical for very long lines. Hence single-mode glass fibers, which are used for long distances, are almost always used with laser sources, and many multi-mode lines are also.

Semiconductor diode lasers, however, do not last as long and are more fragile. Whereas an LED will output light over a fairly high range of currents, the current into a laser diode must be much more carefully controlled. Too little current, and the laser behaves more like an LED — too little light, and not sufficiently pure. Too much current, and the laser burns out. The laser diode also heats up much more than an LED, and its current requirements depend on its temperature. Hence diode laser assemblies often contain a second diode, which is mounted near the laser, and which detects its light. This diode detector is used in a feedback circuit to control the amount of current fed to the laser diode, which complicates the drive circuitry.

Fig. 5-15 shows two laser diodes, which differ only in their mechanical details. The smaller diode on the left needs a separate diode housing which would properly align and attach the laser to the fiber. The larger unit on the right comes with a short length of optical fiber called a *pigtail*. It is already aligned and attached to the diode body by the manufacturer to provide maximum light transfer; you would simply splice the end of the pigtail to the rest of your fiber.

The connection between a laser and the fiber is inherently lossy. For example, the lasers in Fig. 5-15, both made by M/A-COM, are identical GaAlAs diodes (made of gallium, aluminum, and arsenic) which typically output 7 milliwatts of light power at approximately 830 nm in the infrared range. The pigtail is an all-glass 50/125 graded index cable with a numerical aperture of 0.2, whose output is typically only 2 mw. Thus there is over 5 dB loss in the connection.

Before we leave laser diode sources, it is a good idea to just mention safety. Although power levels of 2 or 7 milliwatts seem trivial, they can nevertheless be extremely dangerous because the beam is concentrated into a very small area. If it strikes your hand, you will not even feel it. But if it should enter your eye and be focused by the lens in your eye onto the back of your eye, it can burn the retina and/or optic nerve and cause permanent blindness. Infrared lasers are especially dangerous because you cannot see their beam and may not even be aware that it is on. The standard warning on many laser devices is this: “Do not stare directly into the device or view an operating laser at close range. If viewing is required, the beam should only be observed by reflection from a matte surface utilizing an image converter or by use of a suitable fluorescent screen.”

Radio Shack sells an inexpensive fluorescent screen that can be used to detect infrared light. CCD image sensors of the type used in camcorders are also sensitive to infrared, and can be used to view where a beam is going. But don’t shine the laser beam directly into the camcorder lens, or you will burn the CCD sensor. In fact, when optical fibers are used in short lengths, attenuators in the form of dark filters have to be added to the line to prevent the detector from being damaged.

Light Detectors

At the far end of the fiber, a detector is used to sense the light and convert it back into an electrical signal.

The simplest detector is simply a diode. When a diode is reverse-biased, only a small amount of leakage current flows through it, but when exposed to light, the leakage current increases. This increase in current can then be amplified. An ordinary diode, though, is much too inefficient, especially at the low light levels coming out of a fiber. The light sensitivity can be increased by modifying the diode.

One common photodiode used is the *pin* diode. As the name indicates, the diode consists of three layers — a P layer, an Intrinsic semiconductor layer, and an N layer. By properly doping the layers, the sensitivity can be significantly increased. An even more sensitive diode is the *avalanche* diode, which relies on the avalanche effect to internally amplify its response to light.

It is also possible to use a phototransistor; this is simply a small transistor in a transparent case. Normally, a transistor requires a base current to conduct; otherwise there is just a small amount of leakage current in the collector circuit. But exposure to light produces charges in the base, which then lets the transistor conduct just as though there had been a base current. Phototransistors are quite sensitive, but not fast enough for high bandwidth use.

Physically, detectors are packaged in much the same packages as the LEDs and laser diodes intended for fiber optics, both with and without pigtails. Some detectors also come with built-in integrated circuit amplifiers which simplify the external circuitry.

OTDR

The OTDR or Optical Time Domain Reflectometer is an interesting optical fiber troubleshooting tool, worth mentioning here even though it is far too expensive for many users.

In many ways, the OTDR is similar to the TDR methods we described in Chapter 4. The TDR sends an electrical pulse down a wire transmission line, and observes any reflections which come back from bad terminations or discontinuities in the wire. Similarly, the OTDR sends an optical pulse of light down an optical fiber, and observes any reflections which come back.

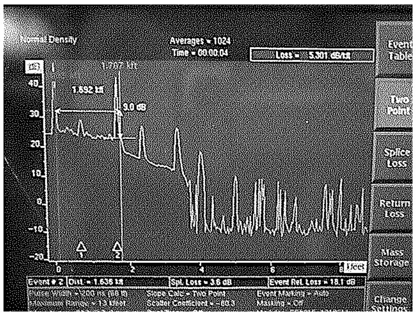

Averages = 1024
Time = 00:00:04
Loss = 5.301 dB/kft

Event # 2
Dist. = 1.636 kft
Spl. Loss = 3.6 dB
Event Rel. Loss = 18.1 dB
Pulse Width = 200 ns (68 ft)
Slope Calc = Two Point
Event Marking = Auto
Maximum Range = 13 kfeet
Scatter Coefficient = -60.3
Masking = Off

Fig. 5-16. Typical OTDR display screen

One difference is that the OTDR is much more sensitive. Not only does it see reflections from bad terminations or connectors, but it also sees reflections from the normal impurities in the fiber. Thus it sees not just bad fiber, but also good fiber.

Fig. 5-16 is a typical display, taken from an OTDR made by Tektronix, and Fig. 5-17 is an expanded view of the important part of the screen. The horizontal axis is the time; since we know the speed at which light travels through the fiber, this axis also shows the distance down the fiber. The vertical axis is the strength of the reflected signal.

The general shape of the curve starts at the upper left (which is close to the instrument, and therefore shows strong signals) and gradually extends to the bottom and right (which is far away, and therefore reflects weaker signals.) The general downward slope shows the attenuation of the fiber.

Periodic tall peaks show very strong reflections from problems in the fiber. These are usually bad splices, sharp bends, or other discontinuities. The actual OTDR instrument is calibrated in time and distance, so you can read out the distance of these discontinuities from the beginning of the fiber. This is of tremendous help in isolating problems.

Conclusion

What kind of information is sent through fibers?

The answer is: everything. Most fiber optics applications today are for digital data. Even when the information is analog, in most cases it is converted to digital data and then transmitted. There are a few instances

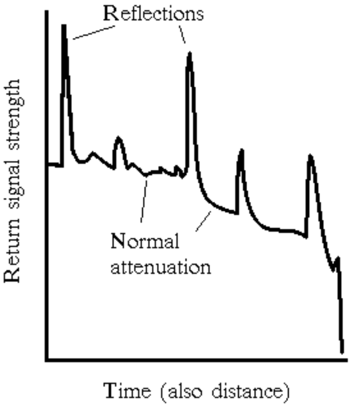

Fig. 5-17. Expanded view of OTDR display trace

where analog data is sent through the fiber by gradually varying the intensity of the light beam — cable TV is a prime example — but even this is likely to change to digital transmission. Digital TV is, after all, just around the corner.

When discussing fiber optic cables, we have also assumed that the fibers are straight, so that once a ray enters a certain mode, it stays there. In reality, fibers bend around corners and obstructions; cables hung on telephone poles will hang down between supports, and then have kinks at the point where they are attached to the pole. Even fibers that look straight have deflections that are large when compared with the wavelength of light. Light rays therefore tend to change modes inside a fiber all the time, and this does increase the attenuation as well as increase the dispersion.

One thing is certain — the last ten years have seen great strides in optical fibers. The quality has shot up, and prices have dropped. The result is that, for many applications, optical fibers are much preferable to copper wire.