2

Why Do We Need Optics?

2.1 The quest for bandwidth

In Chapter 1, the relationship between information content and bandwidth was explored. We saw that when information is cast into the form of a voltage signal, for example, the bandwidth occupied by the signal is greater, the greater the amount of information it represents. When this information is impressed (modulated) on to a carrier in order to send many signals over a trunk telecommunications line, the line itself has to have the capability for a large bandwidth-distance product (that is, the ability to carry a large amount of information over a great distance) if the line is to be economic as a trunk carrier.

A pair of copper wires becomes uneconomic for more than about 300 speech signals over 10 km of distance. With each speech-modulated carrier occupying about 6 kHz of bandwidth, this gives a total bandwidth of about 2 MHz and a bandwidth-distance product of about 20 MHz/km. Clearly, the conduction losses severely limit the bandwidth capabilities of copper wires and hence, for more capacity than this, it is necessary to look away from copper wires. We are looking for means of telecommunication on carrier waves at much higher frequencies, which can travel large distances with little attenuation. Around the time that this became desirable for the further development of telecommunications technology, radio waves were discovered (by Heinrich Hertz, in 1888). These waves were a particular example of a general class of waves known as electromagnetic waves. Before being able to appreciate how these have helped in the advance of telecommunications technology, it will be necessary to take a diversion to understand their primary features.

2.2 Electromagnetic waves

It had become clear during the early nineteenth century that electricity and magnetism were intimately related. Electric currents, moving electrons, were known to give rise to magnetic force fields, and moving magnetic force fields were known to give rise to electric currents.

In 1864 James Clerk Maxwell, a Scottish mathematical physicist, showed that this meant that electric and magnetic force fields ought to be able to reproduce each other continuously in space and thus give rise to a traveling wave: the electromagnetic wave. Maxwell went further and calculated the speed of these waves. He found this to be very close to the known value of the speed of light. It therefore became clear that light itself was an electromagnetic wave. Further experimental evidence for the nature of such waves came 24 years later, when Heinrich Hertz succeeded in producing radiating electromagnetic waves with a varying electric current in a spark chamber. Hertz was able to detect these waves by causing them to produce another spark at an unconnected point a few meters away. He called these waves radio waves, and radio was born (Figure 2.1).

The electromagnetic wave, consisting as it does of mutually sustaining electric and magnetic force fields in space, is depicted in Figure 2.2a. This wave is characterized similarly to that of the pure wave pictured in Figure 1.8a, with an additional feature that results from the fact that it now travels in distance as well as in time. If the wave is observed at one point in space, then the peaks and troughs of either field come and go with time and, as we saw in Section 1.5, the number of peaks (or troughs) that occur in one second is called the frequency of the wave; this is measured in cycles per second (Figure 2.2b), a unit that, as has already been mentioned, is given the name Hertz (abbreviated Hz). Normally, for radio waves, the symbol f is assigned to frequency.

If, now, the wave of just the electric field, say, is observed at one point in time (Figure 2.2c), we see the wave strung out in space, so that the

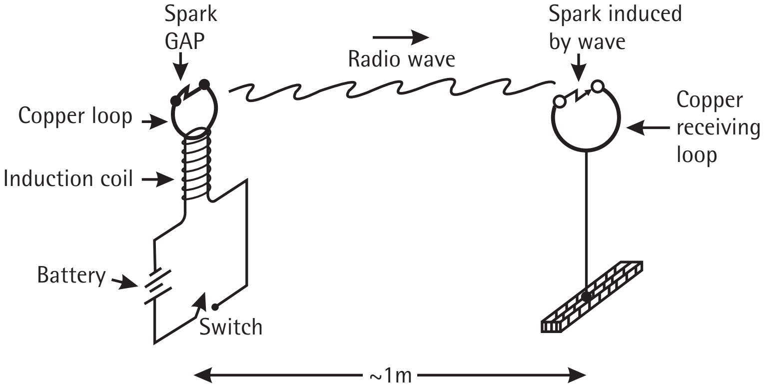

Figure 2.1 Heinrich Hertz's experiment proving the existence of radio waves (1888).

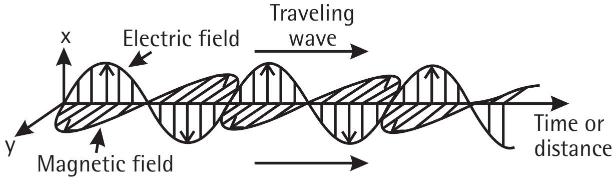

(a) Mutually sustaining electric and magnetic field variations of the electromagnetic wave

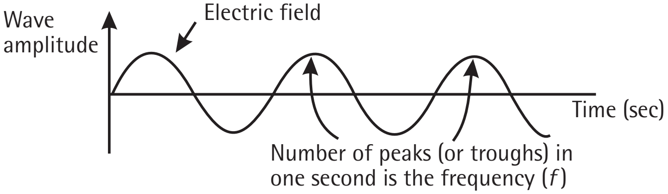

(b) The variation in time at one point in space through which the wave passes

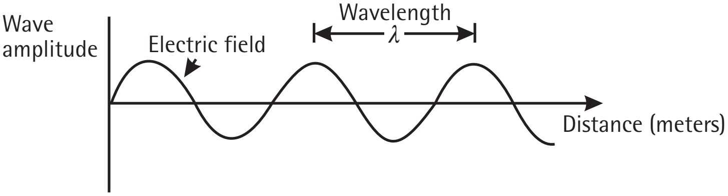

(c) The variation in space at one point in time for the traveling wave

Figure 2.2 Electromagnetic wave definitions.

wave is now also characterized by the distance between peaks (or between troughs). This distance is known as the wavelength, and it is normally assigned the Greek letter λ (lambda). The wavelength is measured in meters or, as it gets smaller, in millimeters (one-thousandth of a meter, 10^{-3} m, written as mm), micrometers (one-millionth, 10^{-6} m, μm), or nanometers (one-billionth, 10^{-9} m, nm). The wavelength comprises an important behavioral feature of the waves, as we shall soon see. It is the same for both the electric and magnetic fields, as is the frequency; in order mutually to sustain each other, the fields must vary in the same way in both space and time. Now, at one point in space through which the wave is traveling there will be f wavelengths passing in any one second, so that the speed of the wave will be $f \times \lambda$ meters per second. This is conveniently written:

$$c = f \times \lambda$$

where c is the speed of the wave. The speed of electromagnetic waves in free space does not depend on their frequency or their wavelength. It is a fundamental constant of nature. Nothing can travel faster. Its measured value is 299,792,458 meters per second and, for most telecommunications purposes, this can be well approximated by a value of 300,000,000 meters per second, or more conveniently in scientific notation, 3×10^8 m/s^{-1}. It travels only very slightly slower in air (the difference is only 0.03%), because the air offers very little resistance to the passage of the waves. Since $f \times \lambda$ is constant, it follows that as the frequency (f) rises, the wavelength (λ) must get smaller; there is no limit to the frequency (or the wavelength) that these waves can have. Hence, for example, a wave at a frequency of 3 kHz will have a wavelength of 100 km (i.e., $3 \times 10^3 \times 10^5 = 3 \times 10^8$), whereas a light wave will have a frequency of around 3×10^{14} Hz and a wavelength of around one-millionth of a meter (1 μm).

The full range of what is called the electromagnetic spectrum is shown in Figure 2.3. It extends from the low-frequency radio waves produced by Hertz's spark, for example, to the very-high-frequency gamma waves produced by some of the most violent objects and processes in the universe, such as black holes and the collapse of giant stars. With such a wide range of carrier frequencies available, it is clear that this spectrum might be of great value to us in telecommunications technology. Let us now look at how these waves can, in fact, be used in pursuit of important advances in telecommunications.

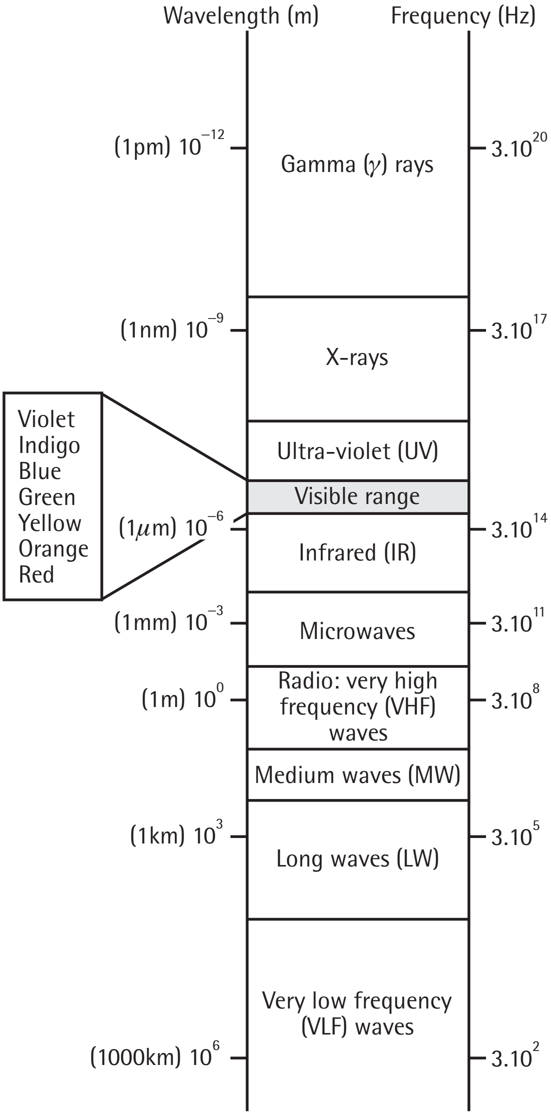

Figure 2.3 The electromagnetic spectrum.

2.3 Transmission with electromagnetic waves

Having learned something about electromagnetic waves, we next ask: What are their advantages for telecommunications? There are several.

The first is that, since the waves travel freely through air and space, no physical connection is necessary between transmitter and receiver (as Hertz discovered). This has enormous advantages, because it means that we will not be hindered by the type of resistance that appears in copper wires, notably at higher frequencies. In fact, Hertz's discovery of the electromagnetic waves known as radio waves was the first great leap forward from copper-wire-based telephony. However, it wasn't until 1894, six years after Hertz's original experiment, that Guglielmo Marconi first succeeded in transmitting information using radio waves.

The second advantage is that, as we have seen, there is now, in principle, an infinite range of frequencies available, and therefore, potentially, infinite signal bandwidth. The real position is not nearly as rosy as this, as we shall soon discover, but nevertheless, electromagnetic waves do offer the possibility for a vast increase in available bandwidth when compared with copper wires.

A third great advantage of using electromagnetic waves is that they travel very fast; we have already noted that their speed is very close to 300,000,000 meters per second; nothing can travel faster than this (this is one of the cornerstones of Einstein's special theory of relativity). It means, of course, that there is very little delay in receiving the transmitted messages: electromagnetic waves can travel all the way around the circumference of the earth in only about one-seventh of a second.

How, then, should we use these waves in our trunk telecommunications channels? To answer this, let us first examine how such waves can be generated. We need to generate waves of electric and magnetic force fields, and this can be done conveniently by forcing electrons to move rapidly back and forth along conducting wires, since pressure of electrons, when they are all compressed at one end of a wire, creates an electric force, and moving electrons, an electric current, creates a magnetic force. These force fields move away from the wires as a pair of mutually sustaining waves to form the electromagnetic wave. These generating wires are then called "antennas." Hertz's spark was effectively an antenna, since it involved rapidly oscillating electrons (albeit now not in a wire), as do all sparks. The problem for telecommunications is that such antennas radiate waves in all directions, rather than just toward a particular receiver along a trunk line (Figure 2.4). This means that any one receiver will receive only a small fraction of the transmitted power, and the farther away it is, the less power it will receive. This wave spreading thus comprises, effectively, a source of attenuation. Of course, it is sometimes very convenient for a transmitter to

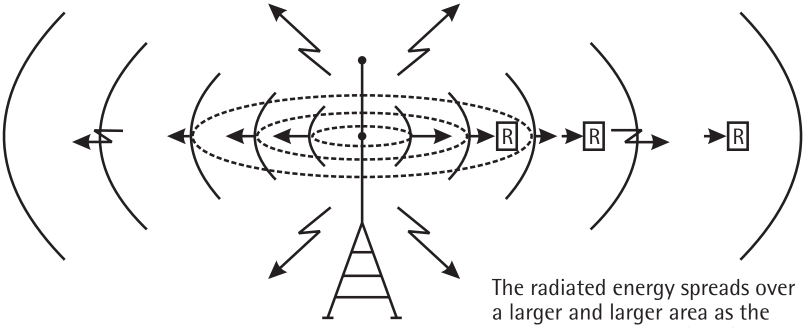

The radiated energy spreads over a larger and larger area as the wave progresses, so that the energy available to a given receiver diminishes with distance

Figure 2.4 Uniformly-radiating antenna.

be received by a large number of receivers spread over a wide area, and this is what happens for national radio and television broadcasts; it is also convenient when the receiver is moving around an area, as in mobile telephony. But these are relatively narrow-bandwidth applications. The rates of information transmission are quite small. For very-wide-band trunk applications, point-to-point links are required. So how can this spreading of the transmitted power be overcome?

One solution is to guide the waves in some kind of conducting arrangement such as is shown in Figure 2.5. Here there is a central conducting wire of copper lying along the axis of an outer copper cylinder, with the space between the two elements filled with an insulator, such as polythene. This is known as a coaxial cable (since both wire and cylinder share the same long axis) and it comprises a flexible guide for electromagnetic waves.

Since conductors are involved in this arrangement, some of the wave power depends on the movement of electrons in the conductors, and it may seem that we are saddling ourselves with the same attenuation problems as in the case of copper wires. However, in this case, most of the power is carried in the electromagnetic wave between the conductors, and only a relatively small amount in the flow of electrons (Figure 2.5). The result of this is that very much higher wave frequencies can be used before the attenuation becomes too severe, and the closed structure prevents the copper wire from itself becoming an antenna, which would cause it to radiate signal power away into the surrounding space (this feature was mentioned as a source of attenuation in Section 1.9). Taking a minimum economic

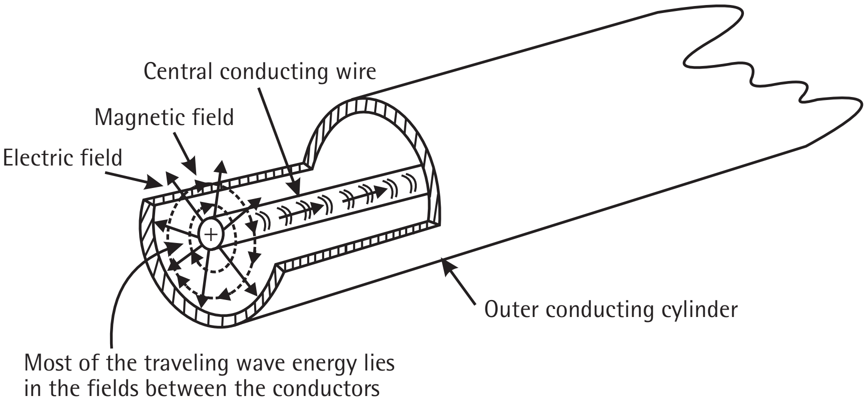

Figure 2.5 The coaxial cable.

repeater spacing of about 20 km, we find that it is possible to transmit around 3,000 speech signals simultaneously over a coaxial cable, corresponding to a bandwidth of about 20 MHz and a bandwidth-distance product of 400 MHz/km. This is about 20 times more than was possible with a pair of copper wires.

Much more bandwidth is still needed, however. So far we have been concentrating on voice channels, but there is much more to telecommunications than just these. There is facsimile (still pictures), video (moving pictures), and computer data, for example. For a video signal (e.g., television) about 5 MHz of bandwidth is required, and for computer data just about as much as possible, in the longer term, anyway. Hence we must look toward higher and higher carrier frequencies. It is clear that the electrical resistance of copper (or any conductor) will eventually always impose a severe limitation on bandwidth, whether it is carrying a current or guiding a wave, so we must look beyond that. The obvious first direction in which to look is toward the free-path transmission of high-frequency electromagnetic waves through the atmosphere. There are three questions that arise immediately in regard to this: First, how can such waves be directed in narrow beams so that point-to-point links can be constructed for trunk lines? Second, how well do such waves travel through the atmosphere? And third, how easy is it to generate such waves? Let us look at these three questions in turn.

First, how can we produce narrow, directed beams of electromagnetic waves? Well, we saw earlier that one way of producing electromagnetic waves is to cause electrons to oscillate rapidly back and forth in a conducting wire known as an antenna. The ideal length of wire needed to generate waves of a particular frequency is equal to about half a wavelength at this frequency. The wire is said to resonate, like a plucked violin string, at this frequency when it has the correct length (Figure 2.6b), just as the violin string will generate sound waves with a half-wavelength equal to its length (Figure 2.6a).

If we are going to produce narrow beams of such waves, it is necessary for us to design structures that are geometrical arrays of these half-wave antennas (known as dipoles) with half-wavelength spacings, so that the waves from the various antennas can reinforce in the required direction and cancel in other directions, by a process known as wave interference. (We will look more closely at this process in Chapter 3). Constructing such arrays is much easier when the wavelengths are small and thus when the antennas are short. For example, medium-wave radio broadcasts, at frequencies of around 1 MHz, operate at wavelengths of about 300 meters (i.e., 3×10^8 m/s^{-1} ÷ 10^6 (Hz) = 300 m), so that very tall masts are required, of order 150 meters high (which is half this wavelength). Clearly, arrays of masts with these heights and spacings are quite impractical. It is much easier to direct the waves when the wavelengths come down to about 1 meter. This corresponds to a frequency of 300 MHz, entering what is known as the

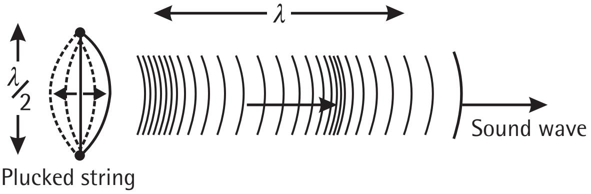

(a) Oscillating string

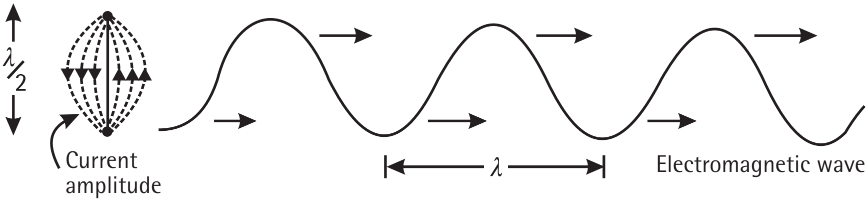

(b) Radiating half-wave dipole

Figure 2.6 The half-wave radiator.

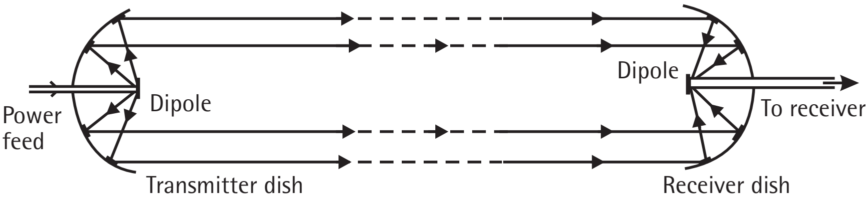

(a) Narrow beam communication with microwave dishes

(b) A microwave dish for satellite reception

Figure 2.7 Microwave dishes. (BT Corporate Picture Library: A BT Photograph.)

"microwave" range (see Figure 2.2a). The added advantage of moving to the higher frequency is, of course, that it also offers greater bandwidth.

Microwave antennas can be quite compact and, clearly, will become more so as the frequency rises. The most effective kind of microwave antenna is the well-known microwave "dish," such as is used for receiving satellite television. This has a diameter of about a meter. The dish is made of conducting material, and the wavelength of the radiation is smaller than the dish diameter. A dipole is placed at the focus of the dish (see Figure 2.7a) and, when fed with electrical power, it radiates in all directions. Wherever radiated waves from the dipole strike the conducting dish, they effectively create secondary, radiating dipoles. Thus all of these dipoles, on the inner surface of the dish, form the equivalent of an ordered array that, owing to the phases of the waves resulting from the geometry, all reinforce in just one direction: the forward direction. Hence the radiation is narrowly collimated in this direction, toward the receiving dish (Figure 2.7a). When it arrives at the receiving dish, the same process happens in reverse, and the incoming radiation is focused on to a dipole placed at this dish's focus. This results in an electric current that then passes to the receiving electronics. A photograph of a larger microwave dish for receiving satellite transmissions is shown in Figure 2.7b.

With such dishes, microwave beams can be transmitted effectively over distances of up to 250 km before beam spreading takes over again and significantly attenuates the signal at the receiver. So now we have the answer to the first question: that of how to produce narrow, directed beams of radiation.

The second question concerns how easy it is to generate such waves. We know that we must make electrons oscillate in a conductor, but to do this at these very high frequencies means that the electrons have to move very fast and hence have to be fed with a lot of energy, and in such a way as to force them to oscillate back and forth at these very high rates. How can this be done?

The problem was first solved satisfactorily during World War II, when the necessity for effective radar systems at microwave frequencies motivated a lot of effort to develop high-power sources of microwaves. Radar played an important part in the defenses of southern England during the Battle of Britain in 1940, and its success was largely due to microwave sources known as the klystron and the magnetron. These devices provided sources of high power at wavelengths of about 10 cm (corresponding to frequencies of about 3,000 MHz). The magnetron is also used today in microwave ovens. There have been other important developments since.

So we now have both the sources and the directional arrays necessary for trunk-line, free-path transmission at frequencies up to 3,000 MHz and beyond, giving us the potential for another large increase in trunk bandwidth-distance. Unfortunately, however, not all of this 3,000 MHz carrier is available as signal bandwidth. The main reason for this is that any given microwave source cannot provide power over all the frequencies up to 3,000 MHz. It is essentially a high-frequency device, and it provides energy within a fairly narrow range of frequencies, of order 5% of the central carrier. This gives a bandwidth of around 150 MHz at 3,000 MHz. At first sight this looks little better than the coaxial cable, which had a bandwidth of about 30 MHz, but only over a distance of about 20 km. Our new microwave bandwidth of 150 MHz can travel a distance of around

250 km in the atmosphere. This gives it a bandwidth-distance product of 37,500 MHz/km, an increase of almost 100 times over that for a coaxial cable.

It is clear that the frequency units have again become cumbersome, so that yet another adjustment is needed. It is convenient now to define a frequency unit equal to 1,000 MHz: the gigahertz (10^9 Hz), abbreviated GHz. The bandwidth-distance product above now becomes 37.5 GHz/km.

Still more bandwidth-distance is needed, so higher and higher we go. There are sources available at 10 GHz (10^{10} Hz) and thus the bandwidth available from these is of order 500 MHz. This is enough to handle 100,000 speech signals or 100 video signals, and again the radiation can be beamed over 250 km or so with the aid of a microwave dish.

Clearly, these beamed microwave signals need a direct line of sight between transmitter and receiver. Microwaves cannot bend around the earth or around buildings, as the much longer wavelengths can do. Roughly speaking, waves can only bend around objects that have a size smaller than their wavelength. A wave can, effectively, curl itself around an object smaller than the rate at which it, itself, is changing in space. At 1 GHz the wavelength is only 30 cm, so that most man-made objects in towns or cities, being larger than this, will block them. It is for this reason that many larger cities have now constructed tall telecommunications towers, such as the BT Tower in London (Figure 2.8). These towers are high, slim structures that support sets of microwave dishes close to their tops, in order to allow direct line of sight transmission paths over the tops of the city buildings to, perhaps, another such tower in another city.

Satellite communications also use frequencies in this range (1 GHz to 10 GHz). An artificial satellite is placed in an orbit where it revolves at the same rate as the earth, a so-called geostationary orbit (that is, at a height of about 35,000 km). It thus remains fixed with respect to the earth's surface and it can be used to "bounce" microwaves between any two points on the earth to which it is visible (Figure 2.9). Clearly, as the waves are going almost straight up and down, there is little possibility of obstruction. One difficulty, however, is that, because of the large return distance, there is now an appreciable time delay, of about one-third of a second, between transmission and reception. This can be troublesome, and irritating, during a two-way conversation.

So now we have a formidable range of carrier frequencies for wideband telecommunications. These are summarized in Table 2.1. They range from the very low frequencies (VLF) carried by copper wires through the

Figure 2.8 The BT telecommunications tower in London. (BT Corporate Picture Library: A BT Photograph.)

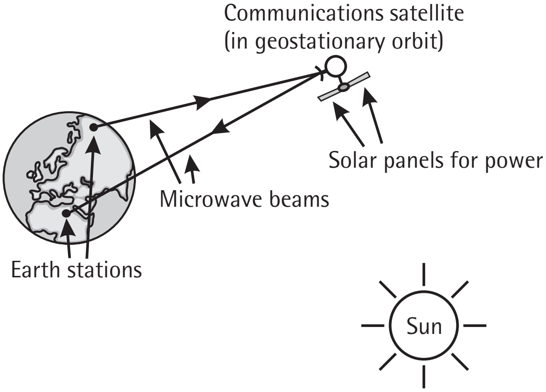

Figure 2.9 Satellite communication.

Table 2.1 The Usage of the Radio Spectrum

Name of Band	Frequency Range	Usage
Very low frequency (VLF)	3–30 kHz	Very long distance communications
Low frequency (LF)	30–300 kHz	National broadcasting Radio navigation
Medium frequency (MF)	0.3–3 MHz	National broadcasting
High frequency (HF)	3–30 MHz	Radio telephony
Very high frequency (VHF)	30–300 MHz	FM broadcasting Television broadcasting Mobile radio telephony Radio navigation
Ultra-high frequency (UHF)	0.3–3 GHz	Television broadcasting Mobile radio telephony Radio navigation Radar
Super-high frequency (SHF)	3–30 GHz	Multi-channel trunk telephony Radar Satellite communications
Extra-high frequency (EHF)	30–300 GHz	TE_{01} waveguide (?)

low (LF), medium (MF), and high (HF) frequencies that can be carried by coaxial cable, up to the VHF and microwave frequencies, which can only be transmitted effectively through the atmosphere and free space.

But the demand for more and more bandwidth is relentless. Higher and higher we must go in our search to meet the ever increasing demands of the computer age. However, at around 10 GHz we encounter another major obstacle. At this frequency the atmosphere begins to absorb, and therefore attenuate, the microwaves quite significantly, especially when it contains moisture in the form of mist, fog, or rain (see Figure 2.10). The trunk telecommunications network, as presently configured (2000), makes good use of microwaves up to about 10 GHz, but what is to be done to acquire even greater bandwidth in the face of this atmospheric blockage?

During the 1970s there was an attempt to guide waves with frequencies up to about 100 GHz (i.e., wavelengths in the millimeter range) in hollow conducting tubes (no central conductors in this case), known as cylindrical waveguides. The idea here was that these tubes could be pumped free of air and moisture and that, therefore, the attenuation would be greatly reduced

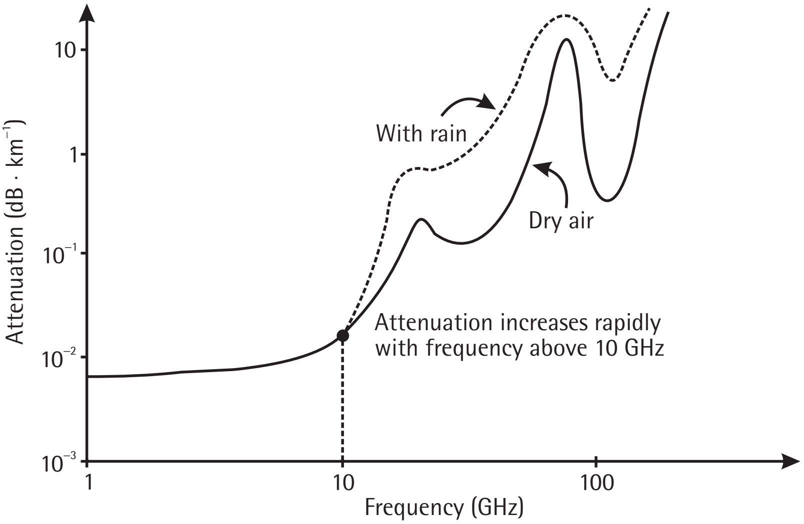

Figure 2.10 The variation of microwave attenuation with frequency for two types of atmospheric condition.

(Figure 2.11). Waveguides for microwaves were not a new idea. They had been used for some time, with either square or rectangular cross-sections, to guide microwaves over short distances (meters), for convenience in taking power to transmitters or from receivers, for example. But the copper conduction losses (the losses discussed in Section 1.7) in them were too severe for use in long-distance transmission at these frequencies. However, the idea that arose in about 1970 was that of using a particular traveling wave pattern, known to the experts as the TE_{01} mode, in which almost all the wave power was carried in the electromagnetic fields, and almost none in the conduction electrons. Consequently, the attenuation was very small indeed; but only when the waveguide was straight. These waveguides were actively researched during the 1970s, but they were temperamental, needed expensive precision engineering, and careful (almost) straight-line installation. Fortunately, other developments were afoot that were soon to overtake the TE_{01} waveguide.

These developments were the result of research, also during the 1970s, on an entirely new type of waveguide. This waveguide contained no metallic conductors at all, and it guided not microwaves but light waves. It was called the optical fiber.

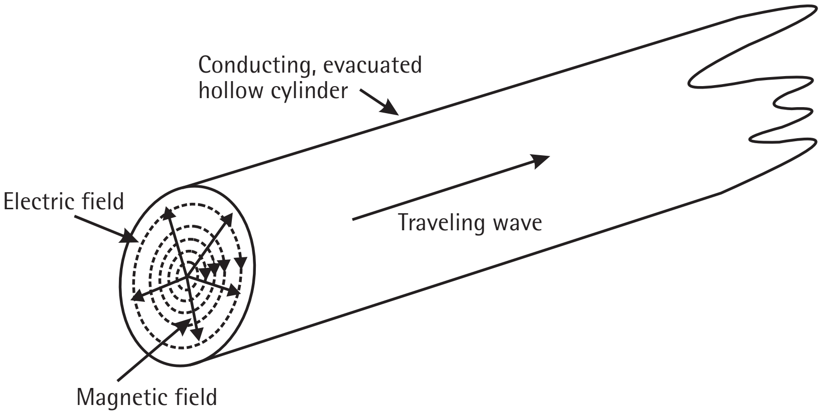

Figure 2.11 High-frequency microwave waveguide (TE_{01} mode).

As we know, light is a form of electromagnetic wave, so the only actual difference between the light waves and microwaves is that of the wave frequency. The optical fiber guides light waves with frequencies of around 300,000 GHz. Clearly, we now need yet another frequency unit: the terahertz (10^{12} Hz), abbreviated THz. So now we are dealing with frequencies of the order of 300 THz, and wavelengths of the order of one-millionth of a meter, or one micrometer (1 μm). This wave frequency lies within what we call the infrared region (see the spectrum, Figure 2.3), which is only just beyond the range of human vision. It is generally classed as "light," since it obeys the laws of optics (which are to be covered in the next chapter) and can, for example, be focused by glass lenses, reflected by glass mirrors, et cetera. The collimation and focusing of microwaves was beginning to show such behavior, as we noted when discussing them (see again, for example, Figure 2.7a). The changes in wave behavior are gradual as we pass along the electromagnetic spectrum and, in any case, are changes of scale rather than of kind.

The range of human vision (i.e., the range to which the human eye's retina is sensitive) is from a wavelength of 0.7 μm, which corresponds to red light, down to 0.4 μm, which corresponds to violet light. Hence, just above 0.7 μm we speak of infrared light and just below 0.4 μm, of ultraviolet light. Between 0.7 μm and 0.4 μm we have the full visual color spectrum: red, orange, yellow, green, blue, indigo, violet. Green light, in the middle of the range, has a wavelength of about 0.5 μm and a frequency of about 600 THz.

For convenience, the notation for all the frequency and wavelength units discussed so far is summarized in Tables 2.2 and 2.3.

Since this new optical waveguide is not made from metal, it involves no conduction electrons and thus no electrical losses. It therefore comprises a very-high-frequency (around 3×10^{14} Hz) carrier system with very low loss. This was just what the telecommunications industry needed. Almost overnight, it rendered the TE_{01} waveguide redundant.

The bandwidth offered by an optical wave, taking 5% of 300 THz, is of order 15 THz, or 15,000 GHz. This is enough for every member of the human race to be talking to another member at the same time—and all along one optical-fiber waveguide. Clearly, this optical regime has the potential to meet all our telecommunications bandwidth requirements well into the foreseeable future, possibly for the whole of the twenty-first century. Additionally, the optical fiber is very thin (about 100 μm in diameter), is of low weight, is easily bent around corners, and is made from sand, one of the cheapest and most abundant materials on our planet. This is why a global optical-fiber telecommunications network is presently under construction.

The economic advantages of being able to pass more and more signal bandwidth down one trunk-line channel are illustrated in Figure 2.12. The

TABLE 2.2 Multiple Units

FREQUENCY (Hz)	SCIENTIFIC NOTATION (Hz)	NAME	ABBREVIATION
1,000	10^3	Kilohertz	kHz
1,000,000	10^6	Megahertz	MHz
1,000,000.000	10^9	Gigahertz	GHz
1,000,000,000,000	10^{12}	Terahertz	THz

TABLE 2.3 SUB Multiple Units

WAVELENGTH (FRACTIONS OF A METER)	SCIENTIFIC NOTATION	NAME	ABBREVIATION
Thousandth	10^{-3}m	millimeter	mm
Millionth	10^{-6} m	micrometer	μm
Thousand millionth	10^{-9} m	nanometer	nm
Million millionth	10^{-12} m	picometer	pm

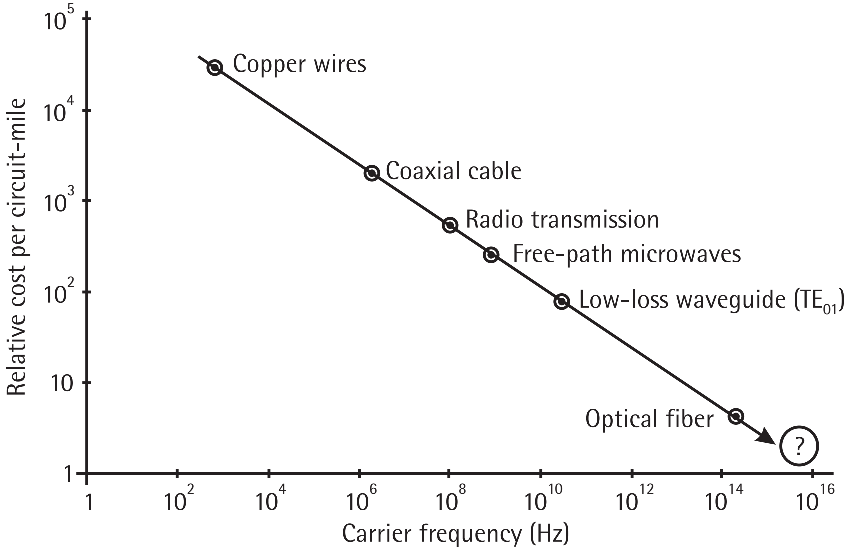

Figure 2.12 The economics of increasing carrier frequency.

diagram shows how dramatically the real (i.e., inflation-adjusted) cost for unit bandwidth-distance product, measured in units of kHz/km, has fallen with increase in carrier frequency, as the changes in technology have occurred. Already, for example, optical-fiber technology has reduced costs by a factor of about 10,000 compared with the early days of copper conductors, and this new technology is only in its infancy. A reduction by another factor of 10 certainly is soon to be achieved and, in the medium term, probably a factor of 100. With such enormous bandwidth capability available at such low cost, there is the near certainty of another major impact on the whole structure of society: on the way we work, play, learn, and generally live our lives.

Now that we have gained an appreciation of the developments that have led to optical-fiber technology and of the position it occupies in the overall scheme of telecommunications, it is time to begin to understand how optical fibers and optical-fiber telecommunications systems actually work. This is the purpose of the rest of the book. Our first task must be to become familiar with the optical fiber itself. This is the subject of the next chapter.

2.4 Summary

In this chapter we have seen how the relentless need to increase transmission signal bandwidth or, more specifically, bandwidth-distance, leads us away from copper wires as the carriers of information, to electromagnetic waves. We studied the basic properties of these waves and then looked at the ways in which these waves can be used in trunk-line telecommunications networks. This led us first to the coaxial cable, on to beamed microwaves through the atmosphere, and then, because of the attenuation by moist air above about 10 GHz, to the evacuated microwave waveguide.

We finally noted that the above systems are being superseded by an entirely new type of waveguide: the optical fiber. This waveguide guides light waves, and it contains no conducting material at all. There are, therefore, no losses of the type caused by electrons flowing in conductors, and the losses are very low indeed in this waveguide. This, coupled with the fact that optical frequencies are very high, at around 300 THz (3×10^{14} Hz), means that optical-fiber technology offers the means for providing all of mankind's telecommunications bandwidth requirements well into the foreseeable future, perhaps for the whole of the twenty-first century.